

Object-Oriented Thinking (OOT)

What Is the “Object,” and What Is Its “Purpose”?

Heather Coffman^a and Derek Cabrera^{b,@}

a: Independent Researcher

b: Cabrera Research Lab

@: dac66@cornell.edu

ABSTRACT: Object-Oriented Thinking (OOT) starts with a simple idea: thinking gets better when we can get our thoughts out of our heads and into a form we can inspect, move around, test, revise, and share. But two questions matter if OOT is to become a construct-valid scientific concept: What exactly is the object? And what is the object for? O-Theory answered both questions in print before the evidence examined here existed: the organization inside any externalized mental model should be representable by Distinctions, Systems, Relationships, and Perspectives (DSRP), and the purpose of making that organization explicit should be to inspect, test, revise, and align the model with reality through the Love-Reality Loop (LRL). Coffman's systematic literature review of 51 publications on structured external representations provides an independent test: the corpus was assembled from adjacent literatures without requiring DSRP, MIO, RIO, or LRL terminology. We applied two source-constrained codebooks to the review's validated analytical records. Among 34 papers containing explicit evidence about internal organization, all 65 identified organizational features were representable through DSRP or combinations of its Patterns; the observed residual rate was zero (0/34), inconsistent with a true residual rate above 8.4% (exact one-sided 95% bound). Among 32 papers containing both an organization claim and a linked purpose claim, no core purpose fell outside direct or enabling LRL functions (0/32; bound 8.9%). The coding did not absorb everything: four papers carried only enabling conditions, and two papers in the wider purpose set fell genuinely outside the LRL. A subsequent replication removed the compression boundary by reading all 51 papers in full — 822,878 words — and extracting 5,735 verbatim statements about what these objects are and what they are for, coded blind by coders who saw no theory, no paper identity, and no indication that anything was under test. Of the statements describing organization, 94.8% are described with the four Patterns; of those describing purpose, 93.4% fall within direct or enabling LRL functions; feedback and fitback co-occur across papers at an odds ratio of 3.38; 87.5% of the 439 individual objects carrying both a composition and a purpose join Pattern-composed structure to a named LRL function; and 46 of 51 papers independently satisfy both conditions. Read semantically across every organizational statement, the literature denies in its own voice that these objects are organized, that they instantiate thought, or that an organized mental model exists at all nine times — 0.38%, in four papers. In plain language: the *object* in OOT is organized thought externalized and explicitly instantiated for the *purpose* of making that thought available for correction, improvement, and alignment. This was the mechanism predicted in advance, with both of its components recovered and linked to each other, from a literature that never used its vocabulary.

KEYWORDS: Object-Oriented Thinking | structured external representations | DSRP | MIO | organization | external cognition | mental models | Love-Reality Loop | model-world alignment | secondary analysis

Cite this paper: Coffman, H., & Cabrera, D. (2026). Object-oriented thinking (OOT): What is the “object,” and what is its “purpose”? *Journal of Systems Thinking*, 6(1), 1–29.

<https://doi.org/10.54120/jost.v26i1.26ootwip6>

1. Introduction

Object-Oriented Thinking (OOT) is a relatively new construct with deep preexisting theoretical [1–8] and explanatory [9–21] lineage. For nearly two decades, work on DSRP Theory, systems mapping, tactile manipulatives, visual mapping, structural cognition, and model-world alignment has been circling the same practical problem: how do we get the organization of thought out of our heads and into a form we can see, touch, inspect, change, share, and test?

The early work was practical before it was formal. DSRP treated thinking as an active process of organizing information through Distinctions, Systems, Relationships, and Perspectives. Thinking at Every Desk [22], ThinkBlocks [11,23,24], tactile cognition [9], and later systems-mapping work [25] pushed that structure into the world: onto paper, whiteboards, blocks, and software. The basic move was simple. When the problem gets too complex to hold comfortably in your head, put the thinking somewhere you can work on it.

The systems-mapping literature made the *object-oriented thinking* character of this work explicit. Castellanos Canales, Lucio Maymon, and Cabrera described effective systems mapping as visual, tactile, object-oriented, social, and embodied [26]. More importantly, they described why. *We externalize cognition when ideas become difficult or complex so that thought can occupy a physically or visually manipulable space.* The medium can change — a napkin, a whiteboard, blocks, Post-it notes, or software — but the underlying problem does not. The representation needs a cognitive architecture capable of carrying the organization of the mental model.

Table 1. Selected historical lineage of Object-Oriented Thinking. Each row names a milestone and its contribution to the through-line that ends at Coffman's independently assembled corpus.

Year	Milestone	Contribution to the OOT lineage
1996–2008	DSRP foundations	Defines thinking as the organization of information through Distinctions, Systems, Relationships, and Perspectives. Identifies the 4 patterns, 8 elements, and 4 dynamics. Establishes theory, existence and effect studies.
2008–2015	Translational, case studies, classroom and tactile work	Moves the theory into visible and manipulable forms; abstract ideas grounded in concrete experience. Early object oriented thinking.
2015 - 2024	Systems Thinking Made Simple; Systems Mapping	Develops DSRP as a practical mapping and thinking architecture. Makes the lineage explicitly

Year	Milestone	Contribution to the OOT lineage
	UCG and Map-Activate-Check; Mathematical theory, Unified Calculus, DSRP-484, MIO/RIO, Oneness, Pareto Moves, and O-Theory	visual, tactile, object-oriented, social, and embodied. Formalizes organization, mental models, real-world organization, and model-world alignment.
2024–2026	O-Therapy, STS standards, OOP/OOT, and FCP	Extends OOT and structural cognition into therapy, professional standards, and software development
2026	Coffman's SLR [27]	Deliberately steps outside the originating OOT literature and assembles an independent adjacent-literature corpus

That lineage matters because OOT is not simply a synonym for “externalizing thinking.” A grocery list externalizes thought. So does a diary, a photograph, or a pile of sticky notes. OOT is concerned with something more specific: making the organization of a mental model explicit enough that the organization itself can be inspected and manipulated. The need for that boundary is larger than it first appears, because in one sense every man-made thing is externalized thought. A sidewalk is a tangible externalization of someone's mental model — “walk here” — that organizes an environment and constrains the thinking of everyone who moves through it. A traditional classroom's layout externalizes someone's model of how learning happens: the teacher stands and talks, the students sit and listen. If externalized thought were the criterion, OOT would be a theory of every artifact. What separates an OOT object from a sidewalk is not whether thought was externalized but what the externalization is for. In an ordinary artifact, the model has been discharged into function: its organization does its work silently, and people comply with it without ever inspecting it. In an OOT object, the organization itself is the point — made explicit so it can be examined, questioned, revised, and reasoned with. For its users, the sidewalk ends deliberation; the representational map exists to continue it.

1.1 From DSRP to DSRP-484

The structural theory underneath OOT has refined its communication over time. DSRP is often summarized as four Patterns — Distinction, System, Relationship, and Perspective — but the DSRP-484 formulation is a much larger claim [2,6,8]. It is not four buckets, four categories, or four generic habits of mind. DSRP is not the dictionary definitions of the four terms— Distinctions, Systems, Relationships, Perspectives. It is a structural calculus [6] made of four Patterns, eight Elements, the special relationships among those Elements, and four structural dynamics: Equality, Coimplication, Simultaneity, and Enclosure [2,6,8].

The eight Elements are identity-other, part-whole, action-reaction, and point-view. *Equality* defines each Pattern through its paired Elements. *Coimplication* means one pole necessarily implies its complement. *Simultaneity* means any element can simultaneously occupy all other structural roles. *Enclosure* means no Pattern exists alone: each Pattern recursively requires the others. That is why DSRP-484 is not a checklist. It is an entangled grammar of organization [8].

$$O \equiv \text{DSRP-484}$$

That shorthand is useful here because this paper is not asking whether the corpus happens to mention four familiar words. It is asking whether the different kinds of *organization* named across an independent literature can be represented by the structural grammar DSRP-484 proposes — or whether the literature forces us to add something else.

1.2 MIO/RIO and the Love-Reality Loop

DSRP's Unified Calculus [6], Law of Structural Cognition [28], and wider O-Theory [2] connect that structural grammar to a larger model of mind and reality. The Law of Structural Cognition states that a mental model is information organized in a particular way:

$$\mathbb{M} = I_{\mathbb{M}} \times O_{\mathbb{M}}$$

The corresponding Law of Structural Reality states that reality also has information and organization:

$$\mathbb{R} = I_{\mathbb{R}} \times O_{\mathbb{R}}$$

MIO and RIO therefore give us two organized structures: the model and the world. The practical problem is that they are not automatically the same. We act from our model of reality, not from reality itself. The question is how to improve the fit between them.

The Love-Reality Loop (LRL) is the formal answer to that problem. It is not merely a list of words like inspect, revise, or improve. Feedback moves from reality toward the mental model:

$$\mathbb{R} \rightarrow \mathbb{M}$$

Fitback moves from the model back toward reality through prediction, decision, intervention, or action:

$$\mathbb{M} \rightarrow \mathbb{R}$$

The ratio at any time can be written as:

$$\frac{\mathbb{M}_t}{\mathbb{R}} = \frac{I_{\mathbb{M},t} \times O_{\mathbb{M},t}}{I_{\mathbb{R}} \times O_{\mathbb{R}}}$$

and the limiting condition is:

$$\lim_{t \rightarrow \infty} \frac{\mathbb{M}_t}{\mathbb{R}} = 1$$

The 1 matters. In the ideal limit, model and reality are in a 1/1 relationship: 1 divided by 1 equals 1. This is not a claim that a human being reaches omniscience. It is a directional claim about alignment. Feedback and fitback provide a recursive mechanism by which the information and organization of the model can be corrected toward the information and organization of reality [2,29,30]. Indeed, that patterned aggregate

trend that differentiates epistemology (mind, or $\mathbb{M}$) from ontology (nature, or $\mathbb{R}$) has been measured as a Cognitive Pareto Law or "Pareto Bias" [5,31,32] and the universality of DSRP-484 has been empirically tested as well [32–36].

1.3 Where OOT sits in the mechanism

This is where Object-Oriented Thinking (OOT) sits in the theory. OOT is the practical move that makes the organized mental model explicit enough to enter the loop. Instead of leaving the model hidden in the thinker, OOT gives its organization an external form — a map, diagram, block, card, whiteboard, software object, or other representation that can be inspected and changed [25].

The “object” is therefore not just a physical object. It is organized thought explicitly instantiated. Systems mapping had already described this in object-oriented, visual, tactile, social, and embodied terms [25]. O-Theory later named OOT directly as a symbolic language for mapping mental models and making DSRP structure visible across scales and contexts [2]. DSRP Therapy uses visible cognition and mapping to work on the structure rather than merely the content of thought [12]. Systems-thinking standards and training materials extend the same architecture into professional practice [16]. Recent programming work applies OOT to the problem of making domain structure explicit before and during software construction, including AI-assisted development [13,15].

So OOT had a history before Coffman’s review. The terminology became sharper over time, but the through-line is stable: organize information, make that organization explicit, and use the explicit model to improve thinking and action.

1.4 Coffman's independent turn

Coffman's systematic literature review (SLR) enters the story here [27]. Its value is precisely that it did not simply review the OOT literature above. Coffman deliberately stepped outside the originating theoretical lineage. The review searched adjacent literatures using visual, physical/tangible, and organizational/verbal forms of external representation rather than requiring or searching for OOT, DSRP, DSRP-484, MIO, RIO, or LRL terminology. The result was a 51-publication corpus assembled from heterogeneous disciplines and vocabularies.

The SLR found seven recurring themes:

1. structured externalization of knowledge;
2. representations as reasoning environments;
3. embodied and material engagement;
4. iterative sensemaking and knowledge construction;
5. shared representations and distributed knowledge work;
6. persistent and reusable knowledge; and
7. human-technology complementarity [27].

Four cross-cutting principles concerned:

- 1. representational fit;
- 2. competence;
- 3. contextual knowledge; and
- 4. integration [27].

The review showed that structured external representations matter. But the synthesis does something more specific than that, and it is visible in its own language. Nearly every load-bearing sentence in Coffman's synthesis has the same two-part shape: information is *organized*, and that organization is explicitly instantiated *so that something can be done with it*. The pattern holds across all seven themes:

Table 2. Paired fragments from Coffman's synthesis of the seven themes [27]. The left column describes what the externalized object *is*; the right column describes what making it explicit is *for*.

The organization itself	What making it explicit enables
"organized external forms"	"available for inspection and further use"
"meaningful organization," as distinct from "simple external storage or display"	"active environments in which reasoning occurs"
"relational structure"	"supporting comparison and inference"; "enabling users to manipulate information"
"structures that can participate in reasoning"	"successive cycles of externalization, inspection, revision, and refinement"
"common reference structures"	"compare perspectives, negotiate meaning, and integrate complementary contributions"
preserved "relationships, provenance, and developmental pathways"	knowledge "revisited, validated, communicated, and potentially reused"
"representational structures and workflows"	differentiated capabilities "interact, inform one another, and permit correction and evaluation"

One sentence of the synthesis carries the whole pattern by itself: *"External representations become consequential not because they exist outside the individual, but because they organize knowledge into structures that can participate in reasoning, support progressive refinement, coordinate distributed activity, preserve developmental context, and enable complementary interaction between humans and technologies"* [27].

But the SLR stopped short of answering two questions that become important once the historical lineage above is put back into view — one question for each column of Table 2.

- 1. First: what does "structure" actually mean? The corpus repeatedly describes information as categorized, decomposed, integrated, connected, mapped, sequenced, constrained, framed, or

viewed. The SLR did not ask whether those apparently different forms of organization — the left column — reduce to a smaller structural grammar.

2. Second: what is the externalized object for? The right column names inspection, reasoning, comparison, revision, coordination, preservation, and reuse, but the SLR did not ask whether those apparently different purposes instantiate a single underlying process.

1.5 The prediction came first

The order of events matters, so we state it plainly. The structural claim — that thinking organizes information through Distinctions, Systems, Relationships, and Perspectives — has been in print since 1996 [37]. The formalization of mental models as organized information (MIO), of reality as organized information (RIO), and of their recursive alignment through the Love-Reality Loop was published in 2024–2025 [2,5,6,38]. Coffman’s corpus was assembled after all of that, by searches that never used any of that vocabulary [27]. The theory did not adapt itself to this evidence. The evidence arrived after the theory, from literatures that had never heard of it.

What the theory predicts is a phenomenon, not a number. O-Theory predicts a sequence:

$$\mathbb{M} = I \times O_{DSRP} \rightarrow E(\mathbb{M}) \rightarrow \mathbb{M} \rightleftarrows \mathbb{R} \rightarrow \text{greater alignment}$$

OOT supplies the explicit object:

$$E(\mathbb{M})$$

and the LRL supplies the recursive dynamic:

$$\mathbb{R} \rightarrow \mathbb{M} \rightarrow \mathbb{R}$$

under the limiting condition:

$$\lim_{t \rightarrow \infty} \frac{\mathbb{M}_t}{\mathbb{R}} = 1$$


If the theory is right, two things should be true of any independently assembled literature on structured external representations. The kinds of organization those papers describe should be representable by DSRP — with nothing left over. And the purposes those papers attach to making organization explicit should be the operations of the Love-Reality Loop, or the conditions that enable those operations — with nothing epistemically central left over. If the theory is wrong, an independent literature is exactly where the leftovers should pile up: structural primitives DSRP cannot represent, and core purposes the LRL cannot absorb.

This paper counts the residuals. We report the observed residual rate in each analysis together with its exact one-sided 95% upper confidence bound — the largest true residual rate that remains consistent with

the data. We set no threshold in advance and defend none after the fact. The prediction is zero residuals; the bound tells every reader how tightly the data pin the truth near zero; and readers who require a specific bar can compare the bound against it themselves.

That is a much stronger test than asking whether people like diagrams or whether external representations help thinking. OOT already had a history, a theory, a formal structure, and a growing applied literature before Coffman assembled the corpus. The value of her SLR is that it gives us a body of literature that did not have to use any of those words. That is what makes the convergence interesting.

Figure 1. Predicted OOT Behavior. DSRP supplies the structural grammar of organization; systems mapping and related tactile/visual practices make organized mental models explicit; MIO and RIO formalize mental and real organization; OOT externalizes the mental model into a manipulable object; and the Love-Reality Loop formalizes recursive model-world alignment under the limit $\lim_{t \rightarrow \infty} M_t/R = 1$. Coffman's SLR provides an independently assembled adjacent-literature corpus against which this preexisting mechanism can be tested.



2. Methods

2.1 Source corpus and evidentiary unit

Both analyses started with the same 51 papers from Coffman's systematic literature review [27]. More specifically, we analyzed the validated Phase 3 analytical records for those papers. The primary text was the Evidence-Based Justification field in the S1 Phase 3 Paper Application worksheet. Paper titles were available as context but were not coded. We also removed the final sentence beginning "This supports..." because that sentence was added during the SLR to assign the paper to one or more themes. We wanted to test the evidence that produced the synthesis, not turn around and use the synthesis as evidence for itself. The paper was our primary unit of analysis. Within each paper, we extracted individual claims and, where necessary, smaller structural or functional features. Those features helped us determine what was present, but the statistical estimates were computed at the paper level. A paper with ten examples therefore did not count ten times more than a paper with one.

There is an important boundary on what this analysis can say. We coded Coffman's validated analytical records, not every sentence of all 51 original papers. These records preserve the evidence used in the SLR, but they are necessarily compressed. When something is not found in a record, we therefore cannot conclude that it is absent from the original paper. We can only conclude that it is not explicit enough in

the analytical record to code. All codebooks, coding records, machine-readable outputs, and statistical scripts are preserved in the replication package.

2.2 The two predictions

The theory makes two predictions about this corpus, and each has a symbolic form and a plain-English form.

Prediction 1 (constitutive). The organization inside any externalized mental model is representable by DSRP:

$$O_{E(\mathbb{M})} \subseteq \langle D, S, R, P \rangle$$

where $\langle D, S, R, P \rangle$ denotes the four Patterns and all their combinations. Equivalently, the residual rate — the proportion of qualifying papers containing at least one organizational primitive that D, S, R, P, and their combinations cannot represent — is zero:

$$\pi_X^{(O)} = 0$$

In plain English: open up any structured external representation in the corpus and ask how its information is organized. You will find Distinctions, Systems, Relationships, Perspectives, or combinations of them — and nothing that forces a fifth structural family.

Prediction 2 (dynamic). The core purposes attached to making that organization explicit are the operations of the Love-Reality Loop, or conditions that enable those operations:

$$F_{E(\mathbb{M})} \subseteq \mathcal{L}_{direct} \cup \mathcal{L}_{enabling}$$

Equivalently, the purpose residual rate is zero:

$$\pi_X^{(F)} = 0$$

In plain English: ask what an explicitly organized representation is for. You will find feedback, fitback, revision, and alignment — or the visibility, persistence, offloading, coordination, computation, and embodiment that make those operations possible — and nothing epistemically central that falls outside.

Measurement. For each analysis we report the observed residual count x out of n qualifying papers, the observed residual rate, and the exact one-sided 95% Clopper-Pearson upper confidence bound:

$$\hat{\pi} = \frac{x}{n}, \quad U_{95} = 1 - \alpha^{1/n} \quad \text{when } x = 0, \alpha = .05$$

In plain English: the prediction is zero leftovers. We report how many residuals we actually found, and the largest true residuals rate the data still allow at 95% confidence.

Table 3. The two predictions and how each was measured. Each prediction is stated in symbolic form alongside its plain-English reading and its measurement. Because both predictions assert zero residuals, evaluation reports observed residual counts and rates with exact one-sided 95% upper confidence bounds rather than threshold-based tests.

Prediction	Symbolic form	Plain English	Measurement
P1: Constitutive	$O_{E(\mathbb{M})} \subseteq \langle D, S, R, P \rangle ;$ $\pi_X^{(O)} = 0$	Every kind of organization inside an externalized representation is representable by D, S, R, P, or combinations — no fifth primitive	Residual count, rate, and exact one-sided 95% upper bound
P2: Dynamic	$F_{E(\mathbb{M})} \subseteq \mathcal{L}_{direct} \cup \mathcal{L}_{enabling} ;$ $\pi_X^{(F)} = 0$	Every core purpose of making organization explicit is a direct LRL operation or an LRL-enabling condition	Residual count, rate, and exact one-sided 95% upper bound
Secondary: Pattern and role structure	—	The four Patterns, four direct roles, and six enabling roles need not occur equally often	Cochran's Q; exact McNemar with Holm correction

2.3 Analysis A: What is the “object”?

Our first analysis asked a simple question: when these papers say that information has been structured or organized, what does that organization actually consist of?

We used the OOT-DSRP Structural Sufficiency Codebook v1.0. Importantly, we did not begin by looking for DSRP. The first pass, called S0, was theory-neutral. It retained only statements that actually told us something about the internal organization of an external representation. Simply calling something “structured,” “organized,” visual, physical, useful, or manipulable was not enough. The text had to tell us how the represented information was organized.

Once a qualifying organizational claim was identified, its structural features were coded as Distinctions (D), Systems (S), Relationships (R), Perspectives (P), combinations of these, unresolved, or a possible additional pattern. We then applied the existing DSRP-v10 codebooks unchanged as a stricter test. Those passes require both poles of a DSRP structure to be explicitly present in the text — for example, both whole and part for a System or both point and view for a Perspective. Finally, we ran an adversarial residual, or X, pass. Any apparent structure that did not initially fit D, S, R, or P had to be tested against each Pattern and combinations of Patterns before it could be retained as evidence of something DSRP was missing.

The primary question was therefore not whether we could find DSRP. It was whether the corpus would force us to add something to it. The X pass is where Prediction 1 was designed to fail.

We also asked whether the four Patterns occurred equally often. Because the same paper could contain more than one Pattern, we used Cochran's Q for the overall matched comparison and exact McNemar

tests with Holm correction for pairwise comparisons. Finally, we repeated the prevalence analysis using only the strict DSRP-v10 results, in which a complete dual-pole DSRP structure had to be explicit in the analytical record. We treat this as a diagnostic rather than a second test of Prediction 1. A complete DSRP structure is especially easy to lose when the evidence from a full paper is compressed into a short analytical record.

2.4 Analysis B: What is its “purpose”?

The second analysis followed directly from the first. Once organized information has been explicitly instantiated, what is it for?

We used the OOT-MIO-LRL Purpose Codebook v1.0. Again, we did not start by looking for the LRL. The first pass, P0, was theory-neutral. It extracted only functions that the analytical record explicitly attributed to an external representation. We did not assume that every benefit of externalization was part of the Love-Reality Loop.

Only after a purpose had been extracted did we ask whether it could be represented by the LRL. Direct LRL functions were coded into four roles: feedback/inspection (FB), fitback/action/testing (FT), revision/learning (RV), and alignment/coherence (AL). We also identified six kinds of conditions that make those operations possible: visibility (VS), cognitive offloading/perceptual access (CO), memory/persistence/traceability (ME), shared coordination/distributed alignment (SC), computational augmentation (CA), and embodied/material access (EB).

That distinction matters. Making a model visible is not itself feedback. Preserving a model is not itself revision. But visibility and persistence can make feedback and revision possible. We therefore kept direct LRL operations separate from LRL-enabling conditions rather than defining everything useful about an external representation as “LRL.”

The coding was deliberately conservative here too. Communication, physicality, engagement, or storage did not automatically qualify as LRL-enabling. The record had to explicitly connect the function to inspection, reasoning, learning, coordination, testing, revision, action, or alignment. Anything that still could not be represented by either a direct LRL operation or an enabling condition was retained as a residual.

For the primary analysis, we narrowed the corpus further. A paper had to contain both an explicit S0 organization claim from Analysis A and an explicit purpose linked to that organized representation. This gave us the cleanest test of the proposed two-part mechanism: information is organized into a model, that organized model is made explicit, and making it explicit allows the model to participate in processes of inspection, testing, revision, action, and alignment.

We also deliberately tried to make the prediction fail in two more ways. First, we recomputed the residuals while treating papers containing only LRL-enabling conditions — and no direct LRL operation — as uncovered. This asks whether the prediction survives a narrower definition of the LRL. Second, we expanded the analysis beyond the organization-purpose pairs to every paper containing a qualifying purpose claim. This asks whether it survives a broader definition of the relevant corpus.

Finally, we tested whether the four direct LRL roles occurred at equal rates, and separately whether the six enabling roles did. Cochran's Q was used for both matched comparisons.

2.5 Replication against full text

Section 2.1 records a boundary on Analyses A and B: they coded validated analytical records, which are compressed, so an unstated structure could not be distinguished from an absent one. We removed that boundary by returning to the source. All 51 papers were obtained in full text — 822,878 words — and read end to end.

Extraction came first and carried no theory. Coders were instructed to copy out, verbatim, every statement a paper makes about what an externalized representation is — its nature, definition, composition, components, structure, or medium — and every statement about what one is for — its purpose, function, benefit, use, or effect. Each statement was tagged as generic (about a class of representation) or specific (about one named artifact) and recorded against the referent it describes, in the paper's own noun phrase. Any statement asserting what a representation is not, does not contain, or is not for was captured and flagged. The instructions prohibited the words DSRP, Distinction, System, Relationship, Perspective, Pattern, LRL, OOT and Cabrera, and prohibited every category label, in every field. Coders transcribed; they did not classify. The harvest is 5,735 statements: 2,372 describing what these objects are and 3,363 describing what they are for.

Coding was blind and came second. Statements were stripped of paper identity and randomized across the corpus, so no coder saw one paper's statements together, and no coder could infer that any hypothesis was under test. Statements describing organization were coded with four neutral labels, presented in non-canonical order — NEST (containment, part and whole), LINK (connection, mapping, dependency), MARK (boundary, category, inclusion versus exclusion), STANCE (vantage, viewpoint, audience) — plus an OTHER category. Statements describing purpose were coded with ten neutral labels plus OTHER. Coders were told that recording OTHER is a correct and expected outcome, never penalized, and were told not to stretch a listed label onto a poor fit. Both codebooks and all thresholds were frozen before production coding.

The predictions were registered in advance with their failure conditions. At least 90% of organizational statements are described using the four; at least 90% of purpose statements fall within a named LRL function; at least 90% of the 51 papers individually clear both an 80% structural bar and an 80% purpose bar; feedback and fitback co-occur across papers above chance; and, as the condition that depends on no category boundary at all, the literature rarely states the contrary of the prediction in its own voice.

Coders were validated before entering the run. An 18-item calibration set contained eight passages built to describe organization the four labels genuinely do not cover — metric distance, a probability distribution, a recursive fixed point, pure cyclic order, a continuous scalar field, a rate, a rank order, and a temporal sequence — six ordinary four-label passages, and four passages asserting a denial. The production coder flagged 8 of 8 planted items, false-flagged 0 of 6 genuine items, and identified 4 of 4 denials. A second

model family flagged 5 of 8 and was excluded under the pre-specified rule. The scheme can return counter-evidence on demand, and the coder that entered production had demonstrated it.

Finally, because the theory holds that the externalized object is an organized mental model made explicit, a denial can be lodged at three levels: that the object has no organization, that it does not instantiate anyone's thinking, or that no organized mental model exists to instantiate. All 2,372 organizational statements were read against those three levels, and each denial found was additionally coded for whether the paper holds it in its own voice or reports it as a position held by others.

The total is 5,735 statements: 2,372 describing what these objects are and 3,363 describing what they are for. Of these, 2,241 and 3,241 respectively carry codable structural or functional content — the remainder name a representation or its medium without stating how it is organized or what it does — and the coverage rates below are computed against those subsets. The denial audit, which asks a semantic question rather than a coding one, is computed against all 2,372 organizational statements.

3. Results

3.1 What does “structured” actually mean?

The first question was straightforward: when these papers say information is structured or organized, what are they actually talking about?

Of the 51 analytical records, 34 (66.7%) contained enough explicit information about the internal organization of an external representation to answer that question. Five contained no qualifying structural claim, and 12 were too compressed to classify reliably. We treated those 12 as unknown, not negative — a caution Section 3.4 vindicates, since all 17 yield organizational statements when the full texts are read. The 34 qualifying records produced 51 structural claims and 65 separately codable organizational features.

All 65 organizational features could be represented through Distinctions, Systems, Relationships, Perspectives, or combinations of them; none were coded unresolved. The adversarial residual (X) pass nominated zero candidates across all 51 records: no feature — including any apparent temporal or dynamic structure — resisted representation by D, S, R, P, or their combinations long enough to be flagged for audit.

Prediction 1 said the residual count would be zero. It was:

$$x = 0, \quad n = 34, \quad \hat{\pi}_X^{(O)} = 0$$

The exact one-sided 95% upper confidence bound was:

$$U_{95} = .084$$

In plain English: we went looking for *organization* that DSRP could not represent, and we found none. And the data do more than report an absence — they bound it. If the true rate of DSRP-exceeding papers in literatures like these were even one in ten, a run of 34 straight zeros would already be improbable; the largest true residual rate still consistent with what we observed is 8.4%. This is a statement about this corpus and these analytical records, not a proof of universal sufficiency. But it is exactly what the theory said we would find, said before this corpus existed, about a literature that never used the theory's vocabulary.

Table 4. Structural results (Analysis A). Qualifying papers contained explicit evidence about the internal organization of an external representation. The residual rate is the proportion of qualifying papers requiring an organizational primitive outside D, S, R, P, or their combinations; strict-v10 counts require both poles of a structure to be explicit in the analytical record.

Quantity	Value
Qualifying papers (of 51)	34 (66.7%)
Structural claims / codable organizational features	51 / 65
Features representable by D, S, R, P, or combinations	65 of 65
Papers requiring a primitive outside DSRP	0 of 34
Observed residual rate	0%
One-sided 95% upper bound on true residual rate	8.4%
Pattern prevalence (family level)	D = 13; S = 20; R = 25; P = 2
Pattern asymmetry	$Q(3) = 32.509, p = 4.09 \times 10^{-7}$
Strict-v10 positive papers	16 of 34 (D = 8; S = 11; R = 1; P = 0)
Strict-v10 asymmetry (diagnostic)	$Q(3) = 19.846, p = .000183$

The four structural families were not equally common: Relationships appeared in 25 qualifying papers, Systems in 20, Distinctions in 13, and Perspectives in only 2. The overall matched comparison rejected equality decisively ($Q(3) = 32.509, p = 4.09 \times 10^{-7}$). Exact pairwise McNemar tests with Holm correction (Appendix Table A1) showed that R occurred more often than D, and that D, S, and R each occurred more often than P. D did not differ significantly from S, and S did not differ significantly from R after correction. So there is a clear asymmetry, but the data support a partial ordering, not a strict D-S-R-P ladder.

The sharpest empirical signal in the structural data is the near-absence of explicit Perspective: two family-level appearances in 34 papers, and zero complete point-view structures under strict v10. This is not an anomaly to be explained away; it is a prior prediction of the theory arriving in a third dataset. The Cognitive Pareto Law or "Pareto Bias" holds that people habitually think from views while almost never marking the point the view is taken from — $P_v \gg P_p$ — a disparity measured across more than three million mapping actions from 34,390 users [5] and independently reproduced by blind coders in the Results sections of 13 scientific disciplines, where Perspective ranked last in 23 of 27 paragraphs [32]. The analytical records here are saturated with views — every named category, component, and dependency is a view from some unmarked point — yet nowhere in the qualifying prose does an author state the point a view is taken from. What the coding recorded, then, is not the absence of Perspective but the predicted absence of the marked point: the same epistemological floor, now observed in a corpus assembled with no knowledge of

the prediction. Nor does this strain DSRP-484's Enclosure dynamic, on which every Pattern recursively requires the others: by Coimplication the point is implicitly present in every view; what the corpus shows is that it is almost never made explicit. Whether the full papers behind these records mark points more often than their compressed summaries do is a question for the full-text follow-on described in Section 4.4.

The strict DSRP-v10 results were much harder to satisfy, as they should have been. Sixteen of the 34 qualifying papers contained at least one complete dual-pole DSRP structure in the analytical record, and the strict counts fell most steeply exactly where the Law of Structural Neglect predicts the differentiated form is least often written down — generic relationships eclipse explicit action–reaction ($R_{gen} \gg R_{ar}$) and views eclipse marked points ($P_v \gg P_p$) [5] – and where dual poles are least likely to survive compression: R fell from 25 family-level papers to 1 strict paper, and P from 2 to 0. A compressed summary might say “hierarchical organization” without naming both the whole and the parts. It may say “dependency” without explicitly stating both ends of the relationship. That is why we kept the two analyses separate. Family-level representability asks: what kind of organization is this? Strict v10 asks: are both structural poles explicitly occupied in the available text? We did not combine the two.

3.2 What is externalized organization for?

The second question was just as simple: once the information has been organized and explicitly instantiated, what is its purpose? What do people do with it?

Forty-seven of the 51 analytical records (92.2%) contained at least one qualifying purpose claim. More importantly for the primary analysis, 32 of the 34 papers containing a qualifying structural claim also contained an explicit purpose linked to that organized representation. Those 32 papers gave us the cleanest test of the mechanism.

Prediction 2 said the residual count would be zero. It was:

$$x = 0, \quad n = 32, \quad \hat{\pi}_X^{(F)} = 0$$

with an exact one-sided 95% upper bound of:

$$U_{95} = .089$$

None of the 32 paired papers required a core purpose outside the combined direct-and-enabling Love-Reality Loop architecture.

Table 5. Purpose results (Analysis B). Paired papers contained both a qualifying organization claim and a purpose linked to that organized representation. A residual is a core purpose representable neither as a direct LRL operation nor as an LRL-enabling condition; the sensitivity rows narrow the definition of the loop (direct-only) or broaden the corpus (all-purpose).

Quantity	Value
Papers with a qualifying purpose claim (of 51)	47 (92.2%)
Paired papers (organization claim + linked purpose)	32 of 34

Quantity	Value
Paired papers requiring a core purpose outside direct + enabling LRL	0 of 32
Observed residual rate (paired)	0%
One-sided 95% upper bound (paired)	8.9%
Paired papers with ≥ 1 direct LRL role	28 of 32
Paired papers with ≥ 1 enabling role	30 of 32
Paired papers with both	26 of 32
Enabling-only papers	4
Direct-only coverage (sensitivity)	4 residuals of 32; rate 12.5%; upper bound 26.4%
Direct-only papers	2
All-purpose corpus (sensitivity)	2 residuals of 47; rate 4.3%; upper bound 12.8%
Direct-role heterogeneity	$Q(3) = 2.700, p = .440$
Enabling-role heterogeneity	$Q(5) = 10.389, p = .065$

But there is an important wrinkle, and it tells us something useful about the mechanism.

Not every paper explicitly described a full feedback-fitback loop. Twenty-eight of the 32 paired papers contained at least one direct LRL role: feedback, fitback, revision, or alignment. Thirty contained at least one enabling role. Twenty-six contained both. Four contained only enabling roles, and two contained direct roles without a separately coded enabling role.

When we deliberately narrowed the definition and treated the four enabling-only papers as uncovered, the residual rate rose to 12.5% with an upper bound of 26.4%. So the prediction does not survive if the LRL is defined by direct operations alone.

That is an important finding, not a failure. It tells us that the literature does not say every external representation is itself a complete Love-Reality Loop. Instead, some representations perform direct work in the loop, while others provide the conditions that allow the loop to operate. Making a model visible, persistent, shareable, computationally tractable, or physically manipulable may not itself be feedback or revision, but it can make feedback and revision possible. The enabling architecture is load-bearing.

The direct LRL roles were spread fairly evenly across the paired papers: feedback in 20, fitback in 17, revision in 14, and alignment in 17. We found no evidence that one direct function dominated the others ($Q(3) = 2.700, p = .440$; pairwise comparisons in Appendix Table A3). Among the enabling roles — visibility 11, offloading 10, persistence 6, shared coordination 13, computational augmentation 10, embodied access 3 — the differences approached but did not reach conventional significance ($Q(5) = 10.389, p = .065$).

Table 6. Direct and enabling LRL-role prevalence among the 32 paired papers. Papers can contain more than one role, so percentages do not sum to 100%. The first four rows are direct LRL operations; the remaining six are conditions that enable those operations.

Code	Role	Papers	Percent
FB	Feedback / inspection / evaluation	20	62.5%
FT	Fitback / action / testing	17	53.1%
RV	Revision / learning / model change	14	43.8%
AL	Alignment / coherence / accuracy	17	53.1%
VS	Visibility / explicit availability	11	34.4%

Code	Role	Papers	Percent
CO	Cognitive offloading / perceptual access	10	31.3%
ME	Memory / persistence / traceability	6	18.8%
SC	Shared coordination / distributed alignment	13	40.6%
CA	Computational augmentation	10	31.3%
EB	Embodied / material access	3	9.4%

The wider 47-paper purpose analysis gave us an important check on whether the codebook was simply swallowing every conceivable purpose into the LRL. It was not. Two papers produced genuine non-LRL residuals — an observed rate of 4.3%, with an upper bound of 12.8%. Both involved material representations of physical-activity data. Their purposes included things such as material reward and motivation, self-expression and identity, personal attachment, decoration, jewelry, gifting, use as a conversation piece, and social signaling.

We did not force those functions into the LRL. The available text did not connect them to inspecting a model, testing it, revising it, learning from it, coordinating around it, or bringing it into better alignment with reality. So they stayed outside. Those two papers also lacked a qualifying S0 claim about the internal organization of the representation, so they never entered the primary paired analysis. But their survival as residuals is useful because it demonstrates that the coding scheme could identify, "No. This purpose is doing something else."

3.3 The two pieces fit together

These two analyses answer different questions. Analysis A asks: what is the OOT — the organization in the externalized object — made of? Analysis B asks: what is that organization's purpose? That distinction matters because the findings line up at two different theoretical levels.

At the first level, the heterogeneous structural language used across the literature — categories, components, hierarchies, mappings, dependencies, connections, views, and so forth — was representable through D, S, R, P, or combinations of them. In the language of MIO: for the qualifying organizational features in this corpus,

$$O_{E(\mathbb{M})} \subseteq \langle D, S, R, P \rangle$$

At the second level, the purposes attached to those organized representations — inspection, comparison, testing, revision, learning, coordination, persistence, action, and alignment — were representable through direct LRL operations or conditions that make those operations possible:

$$F_{E(\mathbb{M})} \subseteq \mathcal{L}_{direct} \cup \mathcal{L}_{enabling}$$

Both components of the predicted mechanism therefore co-occur — with the purpose explicitly linked to the organized representation — in 32 of 34 qualifying papers:

$$\mathbb{M} = I \times O_{DSRP} \rightarrow E(\mathbb{M}) \rightarrow \mathbb{M} \rightleftarrows \mathbb{R} \rightarrow \lim_{t \rightarrow \infty} \frac{\mathbb{M}_t}{\mathbb{R}} = 1$$

In plain English: the “object” in Object-Oriented Thinking is organized thought explicitly instantiated. Its “purpose” is to make that organized thought available to the processes by which models can be inspected, tested, corrected, shared, and improved.

There is an obvious selection boundary here. Coffman’s SLR was designed to find external representations containing represented information and explicit organization. Additionally, papers required the representations to explicitly contribute to understanding, reasoning, learning, communication, or decision-making. So it would be circular to announce that the corpus contains organization or that the representations are useful. But that boundary is not an accident of her method — it is the edge of the construct itself. Because every artifact externalizes someone’s thought (a sidewalk no less than a concept map), some line must separate objects that present their organization for thinking from objects whose organization silently performs a function. The SLR’s inclusion criteria — explicit representation, explicit organization, in service of understanding, reasoning, learning, communication, or decision-making — are that line. The corpus therefore samples precisely the class of objects OOT is about.

But that is not what these analyses asked. The SLR did not select for D, S, R, and P as the particular structure of organization. It did not require feedback, fitback, revision, and alignment as the particular dynamic functions of externalization. It did not require the enabling architecture of visibility, offloading, persistence, coordination, computational augmentation, and embodied access. And it did not require those two levels to appear together in the same papers, with the purpose tied to the organized structure.

Yet that is what emerged. The result is therefore more than thematic similarity. It is mechanistic convergence. The corpus independently reproduces both sides of the proposed mechanism: a structural account of what the object is, and a dynamic account of what making that object explicit is for.

3.4 The same test against full text

Every paper speaks. All 51 papers yielded statements of both kinds. That includes all 17 papers whose analytical records had been recorded as carrying no qualifying claim or as too compressed to classify — one of which produced 41 organizational statements from full text after producing none from its record. The qualifying denominators of 34 and 32 papers in Sections 3.1 and 3.2 were properties of the records, not of the literature. Read in full, the corpus is not sparse. It is dense.

What the object is made of. Of 2,241 statements describing organization, 2,124 — 94.8% — are described using at least one of the four Patterns by coders who had never encountered them. Relationship appears 1,362 times, System 955, Distinction 815, and Perspective 255. Thirty-six of 51 papers exhibit all four, and no paper in the corpus exhibits only one. Where both poles of a structure could be named, Relationship named both in 91% of uses, System 86%, Distinction 73%, and Perspective 51% — Perspective being both the least frequent Pattern and the least often complete, the Pareto asymmetry appearing on two independent measures at once.

What the object is for. Of 3,241 statements describing purpose, 3,028 — 93.4% — fall within a direct or enabling LRL function. Direct functions: revision 740, fitback 428, feedback 284, alignment 92. Enabling functions: visibility 1,157, computation 951, coordination 453, embodiment 334, persistence 333, offloading 302. Fifty of 51 papers carry at least one direct and at least one enabling function, and every paper naming enabling functions also names direct ones. Forty-eight name two or more direct functions and 18 name all four.

The loop closes. Feedback and fitback are predicted to be paired moments of one recursive process rather than two independent benefits, and they behave that way: 36 papers carry both, 4 carry feedback alone, 8 carry fitback alone, and 3 carry neither — an odds ratio of 3.38. Revision accompanies fitback in 42 papers. A list of unrelated benefits would produce no association. This one could have come back flat and did not.

The object and its purpose, joined. Because every statement was recorded against the referent it describes, composition and purpose can be joined on the externalized object itself rather than inferred from co-occurrence within a paper. Four hundred thirty-nine distinct objects carry both a statement of what they are made of and a statement of what they are for. Of those, 384 — 87.5% — are composed of at least one Pattern and serve at least one named LRL function. One hundred forty-five are composed of two or more Patterns and serve a direct LRL function specifically. Fifty of 51 papers contribute at least one such object, the median paper contributing seven.

Paper level. Forty-six of 51 papers — 90% — independently clear both an 80% structural bar and an 80% purpose bar using nothing but the blind coders' own label assignments. Forty-nine clear the structural bar and 48 the purpose bar. The three that fail the purpose bar are all material- or tangible-artifact papers, one of them the paper already placed outside the loop and excluded from the paired set in Section 3.2.

The contrary position, located. Reading all 2,372 organizational statements against the three levels at which the thesis can be denied, 17 denials appear and 9 are held in the paper's own voice: 0.38%, in four papers of 51. Two deny that an object is organized, and neither survives inspection — one is the base rung of a DIKW hierarchy in which each higher level adds organization, and the other observes that mindmaps are "essentially radial networks without link labelling," which asserts Distinction and System while complaining that Relationship is left implicit. Six are a single paper arguing an explicitly anti-representational design position, which rejects artifacts as "tangible information carriers" and replaces them with what it calls a tangible mapping offered by "the total set in its spatial organization." One paper denies an organized mental model in its own voice, offering an account "without appealing to mental representations." The sharpest statement of the opposing view in the whole corpus — that external representations are "nothing but inputs and stimuli to the internal mind" — appears twice, both times reported rather than held, in a paper whose own contribution is externalizing latent structure. The literature names the alternative to this thesis chiefly in order to move past it.

A note on residuals. Beyond the four labels, coders were free to flag additional organizational content and did so on 675 statements. These are overwhelmingly additions rather than alternatives: 579 of the 675 also carry Pattern codes, and 70% name spatial layout, order and time, or magnitude — the axes along

which Patterns are instantiated rather than rival primitives. Assessed against a representability rule frozen in advance, 582 resolve into combinations of the four, 88 prove to be properties of an artifact rather than forms of organization, and 5 remain unresolved. We report this for completeness and do not rest the case on it; Section 4.4 explains why.

4. Discussion

4.1 What Structure Is Made Of

Coffman's SLR showed that structured external representations matter across a surprisingly wide range of fields and uses [27]. What it did not try to determine was what *structure itself* is made of. Across the literature, authors used different disciplinary language—categories, components, hierarchies, mappings, dependencies, connections, coordinated views, and other forms of organization. The secondary analysis asked whether those apparently different structures reduce to a smaller common grammar.

At the aggregate or "family" level, they did. The organizational features identified in the qualifying records could be represented through Distinctions, Systems, Relationships, Perspectives, or combinations of them. The residual audit also allowed for the possibility that the corpus would require something else. It did not. No additional organizational primitive was retained: no fifth structural family, no separate temporal primitive, and no separate dynamic primitive was needed to represent the organizational features identified in the available records.

That result is more specific than saying the papers contained "structure." It suggests that the many different ways these literatures talk about organization may be surface variations of a much smaller underlying grammar. In DSRP-484, that grammar is not merely four categories. Distinctions, Systems, Relationships, and Perspectives are recursively related structures composed of eight co-implicating elements and governed by the dynamics of Equality, Coimplication, Simultaneity, and Enclosure. The present test operates at the Pattern-family level, with the strict-v10 analysis probing the Element level; the four dynamics were not separately coded and remain for the full-text study. Within that scope, the result is consistent with the stronger proposition: apparently heterogeneous forms of organization can be represented through the same underlying structural architecture.

The strict-v10 analysis asks a harder question. Instead of asking whether a structural feature is representable through a DSRP family, it requires both poles of the structure to be recoverable directly from the available text. Those results were much sparser, particularly for Relationships and Perspectives. The two results should therefore not be conflated. Family-level representability concerns the *kind of organization present*; strict recovery concerns how completely that organization is stated in the analytical record.

The useful conclusion is simple: *the corpus does not merely contain organization; the organization it contains repeatedly resolves into the structural grammar predicted by DSRP-484.*

4.2 The purpose side is just as specific

The second result may be even more interesting because it asks why people bother making organized thought external in the first place. The synthesis repeatedly uses a two-part structure that can be paraphrased as: organize information in order to do something with it. What follows that phrase is not random. The papers talk about inspection, comparison, inference, testing, revision, learning, decision support, coordination, persistence, computational augmentation, and alignment.

That pattern looks like the Love-Reality Loop [2,28]. Some of those functions are direct operations in the loop: feedback, fitback, revision, and alignment. Others are the plumbing that lets the loop work: making a model visible, keeping it around long enough to revisit, sharing it with other people, offloading some of the cognitive burden, or allowing computation or physical manipulation to participate.

The negative controls are what make this result worth taking seriously. Four paired papers contained enabling conditions without an explicit direct feedback-fitback-revision-alignment role. When we treated those four as failures, direct-only coverage did not hold. That tells us not to overstate the result. External representations do not all explicitly enact a complete LRL. Some simply create the conditions under which the loop can happen.

More importantly, two papers in the wider purpose set really did fall outside the LRL. Their purposes included reward, decoration, gifting, identity expression, personal attachment, and social signaling. We left them outside. If LRL were just a fancy name for anything useful, those cases would have been easy to force into it. We did not. The result therefore supports a more specific claim: the dominant epistemic and adaptive purposes attached to explicitly organized external models are covered by direct LRL operations or by conditions that explicitly enable those operations.

4.3 The result is the mechanism, not a statistic

The most important feature of these findings is not any single number. It is the specificity of what was recovered, given what was predicted, and when.

A generic theory that says external representations help cognition could easily predict that people report benefits. It does not specifically predict that the organization inside those representations will reduce to D/S/R/P families. It does not predict that the purposes linked to making that organization explicit will reduce to feedback, fitback, revision, alignment, and the infrastructure that makes those operations possible. And it certainly does not predict that both reductions will appear together in the same papers, linked to one another — organization and purpose attached, exactly as the organize–externalize–loop–align mechanism requires.

O-Theory made exactly those predictions, in print, before this corpus existed. Coffman's corpus was not assembled by searching for DSRP or LRL language. Yet the secondary analysis recovered both the proposed constitution of the model and the proposed work that making the model explicit enables — with no residual in either primary analysis, upper bounds under 9%, and a coding scheme that demonstrably knew how to say no. Section 2.5 describes the replication that then put the same predictions to 51 full texts and 5,735 verbatim statements, under coders blind to the theory and to the existence of a hypothesis,

and Section 3.4 reports what it found. That is why we describe the result as mechanistic convergence rather than merely thematic similarity.

4.4 Limitations

The principal boundary on Analyses A and B — that they coded compressed analytical records rather than full text — has been removed by the replication in Section 3.4, which recovers the same result from 822,878 words of source material and restores 17 papers the records had rendered unclassifiable. Two boundaries remain worth stating precisely. The first is what a residual rate can mean. When coders may flag any additional organizational content they do so on 30.1% of statements; when that content is assessed against an explicit rule for what the four Patterns cover, 0.23% remains unresolved. Both figures come from the same statements and the same coders, and the distance between them is a property of the predicate rather than of the literature, which is why the case does not rest on a residual rate in either direction. It rests on what no predicate governs: that 94.8% of organizational description reaches for the four Patterns with no theory in the room, that 90% of papers independently satisfy both bars on blind coding alone, that feedback and fitback co-occur at an odds ratio of 3.38, and that the literature asserts the contrary in its own voice nine times in 2,372 statements. The second boundary is coder independence: the production run used a single model family, the second having failed the pre-specified calibration and been excluded, and each statement was coded once. Flagging rates for additional content varied from 14% to 53% across batches randomized across papers, locating that variation in coder latitude rather than in the corpus. Three papers were obtained with truncated text and their counts are floors. All instruments, the calibration set with its key, every coder's exact input and output, and a script regenerating every figure reported here are in the replication package. As to the original analyses:

The analysis tested Coffman's validated analytical records rather than independently recoding every original full text, and both analyses were a single-coder primary run: no second coder or inter-coder reliability statistic exists for these codes, and the X pass, available at every record, nominated zero candidates — a result whose interpretation (true structural sufficiency versus a permissive predicate) is precisely what independent coders on held-out full text will adjudicate. All of this is preserved in the replication package. Because analytical summaries compress source material, an absent structure or purpose cannot be treated as proof that it was absent from the underlying paper.

There is also an important selection boundary. Coffman's SLR was designed to identify explicit representations, explicit organization and a contribution to understanding, reasoning, learning, communication, or decision-making. We therefore do not treat the mere presence of organization or useful cognitive functions as an independent finding; as Section 3.3 notes, those criteria are the definitional line separating OOT objects from artifacts generally, so their presence in the corpus is a feature of the sampling frame, not a discovery. The result tested here is more specific: whether the *particular forms of organization* described in the corpus resolve into DSRP-484, and whether the *particular purposes* attached to making that organization explicit resolve into direct or enabling LRL functions.

The findings support mechanistic convergence on this corpus.

5. Conclusion

Coffman's SLR established the broad empirical terrain. Across disciplines, external representations become useful when information is organized into forms that people can inspect, reason with, revise, coordinate around, preserve, and act on [27]. This secondary analysis asked the next two questions: what is that organization made of, and what is making it explicit for?

The theory had already answered both, in print, before the corpus existed. The corpus agreed. Within the available analytical records, no qualifying organizational feature required a structural primitive outside DSRP (0 of 34; the true residual rate is bounded below 8.4%). And among papers that contained both an explicit organization claim and a linked purpose claim, no paper required a core purpose outside the combined direct-and-enabling LRL architecture (0 of 32; bounded below 8.9%). Those are two different findings, but they fit together:

$$\boxed{\mathbb{M} = I \times O_{DSRP-484} \rightarrow E(\mathbb{M}) \rightarrow \underbrace{\mathbb{R} \rightarrow \mathbb{M}}_{\text{feedback}} \rightleftarrows \underbrace{\mathbb{M} \rightarrow \mathbb{R}}_{\text{fitback}} \rightarrow \lim_{t \rightarrow \infty} \frac{\mathbb{M}_t}{\mathbb{R}} = 1}$$

The first tells us what the object appears to be made of. The second tells us what making that object explicit allows us to do. Together they recover the two-level mechanism predicted by O-Theory, from a literature that never used its vocabulary.

That gives us a tighter post-SLR definition of Object-Oriented Thinking: Object-Oriented Thinking is the externalization of an organized mental model so that its structure can participate in inspection, comparison, testing, revision, coordination, persistence, action, and model-world alignment. DSRP specifies what the externalized organization consists of. The Love-Reality Loop specifies the principal epistemic and adaptive work that the explicit instantiation of thought enables.

That is a more specific claim than saying external representations help people think. It says what the object is and what the object is for — and it said so before this evidence existed. That replication has now been run, and the mechanism survived it. Reading all 51 papers in full and extracting 5,735 verbatim statements, 94.8% of what the literature says these objects are made of is described with the four Patterns by coders who had never heard of them; 93.4% of what it says they are for falls within direct or enabling LRL functions; feedback and fitback co-occur at an odds ratio of 3.38 rather than appearing as unrelated benefits; 439 individual objects carry both a composition and a purpose, and 87.5% of them join Pattern-composed structure to a named LRL function; and 46 of 51 papers independently satisfy both conditions on blind coding alone. Against all of that, the literature asserts the contrary in its own voice nine times in 2,372 statements, in four papers, and the sharpest version of the opposing view appears only as a position reported and then set aside. A theory that merely said external representations aid cognition would predict none of this. O-Theory predicted all of it, in print, before the corpus existed.

How the operational Definition of OOT changed across analyses.

Operational definition entering the SLR:

OOT = externalizing thinking.

After the SLR:

OOT = thinking with and through organized external representations so that thinking can be inspected, reasoned with, revised, coordinated, and improved.

After the secondary analysis:

OOT = making the invisible organization of a mental model explicit so that it can be inspected, tested, revised, and brought into better alignment with reality. DSRP-484 specifies what that organization is; the Love-Reality Loop specifies what making it explicit is for.

$$M = I \times O_{DSRP-484} \rightarrow E(M) \rightarrow \underbrace{R \rightarrow M}_{\text{feedback}} \rightleftharpoons \underbrace{M \rightarrow R}_{\text{fitback}} \rightarrow \lim_{t \rightarrow \infty} \frac{M_t}{R} = 1$$

In plain English:

OOT makes the invisible structure of thought visible and manipulable. DSRP-484 specifies what that structure is; the LRL specifies what making it explicit is for.

The *Unified Calculus* derived this architecture from the top down, as axioms. This corpus supplies what a derivation cannot: an ecological test against researchers in visualization, education, knowledge management, tangible interaction, robotics, external memory, whiteboarding, and software design describing, in their own terms, what these objects are and what they are for. Across 5,735 coded descriptions, 94.8% of the structural descriptions and 93.4% of the functional purposes were captured by the architecture predicted by DSRP and the Love-Reality Loop. The convergence is therefore not merely deductive. The same architecture appears from the bottom up in a heterogeneous literature produced largely outside the theory itself. Most importantly, O-Theory specified that architecture *before this corpus was analyzed against it*. What was derived axiomatically from the top down is independently visible, at scale, from the bottom up.

References

1. Cabrera D. The grammar of reality. *Journal of Systems Thinking*. 2025;5: 1–40.
2. Cabrera D, Cabrera L. O-Theory: The Science of Organization. *Journal of Systems Thinking*. 2025;5: 1–12.
3. Cabrera D. *Systems Thinking*. Cornell University; 2006.
4. Cabrera D, Cabrera L. The Grammar of Reality: The Organizing Principles of Structure and Emergence. *Journal of Systems Thinking*. 2025;5: 1–18.
5. Cabrera D, Cabrera L. The Pareto structure of thought: empirical discovery of the six foundational mental moves. *J Syst Think*. 2025;2025: 1–41.
6. Cabrera D. The unified calculus of organization, structure, emergence, and meaning: The patterns that connect mind and nature. *Journal of Systems Thinking*. 2025;5: 1–41.
7. Cabrera D, Cabrera L, Lobdell. System Thinking. *Journal of Evaluation and Program Planning*. 2008; 299–310.
8. Cabrera L. CD. DSRP-484: Building an advanced understanding of O-Theory, DSRP theory, and the DSRP calculus. *Journal of Systems Thinking*. [cited 17 Aug 2026]. doi:10.54120/jost.v25i1.dc6O4846
9. Cabrera D, Cabrera L. The World at Our Fingertips: The Sense of Touch Helps Children Ground Abstract Ideas in Concrete Experiences. *Scientific American Mind*. 2010;21: 36–41.
10. Cabrera Research Lab. How to Make Invisible Thinking Visible. 2026. Available: <https://podcasts.apple.com/au/podcast/127-how-to-make-invisible-thinking-visible/id1738320168?i=1000775945779>
11. Systems Thinking Standards Institute. Systems Thinking 3-D Whiteboard Blocks. Available: <https://www.stsi.pro/product-page/systems-thinking-3-d-whiteboard-blocks>
12. Cabrera D, Cabrera L, Silberman D, Vincenzes KA, editors. *O-Therapy Manual: Clinician's Guide to O-Therapy, Standards, and Protocols*.
13. Silberman D, Bhuan SKV, Cabrera D. Object-Oriented Programming Requires Object-Oriented Thinking: From Information to Organization in the Age of AI. *Journal of Systems Thinking*. 2026;6: xx–xx.
14. Cabrera Research Lab. Object-Oriented Thinking Changes Everything. 2026. Available: <https://www.youtube.com/shorts/kzp734K5LQk>
15. Silberman D, Bhuan SKV, Cabrera D. Feature-Controlled Programming: Embedding Structural Commitments in AI-Assisted Software Development: An Industrial Case. *Journal of Systems Thinking*. 2026;6.

16. Systems Thinking Standards Institute. The Global Systems Thinking Standards (STS): Theory, Practice, and Evaluation Guidebook. Version 2.3. Systems Thinking Standards Institute (STSI); 2025 Sept.
17. Solin J. Map-Activate-Check: A Systems Thinking-based Approach to Curriculum and Program Design. *Journal of Systems Thinking*.
18. Cabrera Research Lab. Stop Reading About Systems Thinking---Practice This Instead. 2026. Available: <https://www.youtube.com/watch?v=sRmTcRmxhsQ>
19. Cabrera D. Making Smarter Moves with Systems Thinking. 2025. Available: <https://www.youtube.com/watch?v=lz-b8pu9xT4>
20. Cabrera Research Lab. The Power of Putting Ideas on the Table. 2026. Available: <https://www.youtube.com/watch?v=MPQOHtyKnSI>
21. Cabrera Research Lab. How to Learn When Words Don't Work for You. 2026. Available: <https://www.buzzsprout.com/2279323/episodes/18995895-115-how-to-learn-when-words-don-t-work-for-you>
22. Cabrera D, Cabrera L. *Thinking at Every Desk: Four Simple Skills to Transform Your Classroom*. Norton Books; 2011.
23. Cabrera D, Cabrera L. *ThinkBlocks Metacognitive Tactile Manipulatives*. Ithaca, NY; 2006.
24. Cabrera D, Cabrera L. *ThinkBlock Case Studies*. 2006.
25. Castellanos Canales A, Lucio Maymon P, Cabrera D. The Universal Cognitive Grammar of Systems Mapping: A Rubric to Evaluate the Various Tools and Techniques of Systems Mapping. *Journal of Systems Thinking*.
26. Castellanos Canales A, Lucio Maymon P, Cabrera D. The Universal Cognitive Grammar of Systems Mapping: A Rubric to Evaluate the Various Tools and Techniques of Systems Mapping. In: Cabrera L, Cabrera D, Midgley G, editors. *The Routledge Handbook for Systems Thinking*. London, UK: Routledge Press; 2021.
27. Coffman H. Object-Oriented Thinking (OOT) Beyond Its Own Literature. A Systematic Literature Review. *Journal of Systems Thinking*. 2026.
28. Cabrera D, Cabrera L. The Structure of Thought and Reality: Empirical Evidence for the Three Universal Laws Linking Mind, Matter, and Meaning. *Journal of Systems Thinking*. 2025;5: 1–20.
29. Cabrera D, Cabrera L. *Oneness*. Odyssean Press; 2025.
30. Klein M. The asymptotic boundary of O-Theory: Why $1/1 \neq 1$. *Journal of Systems Thinking*. 2026;6: 1–8.
31. Di Marco G, Cabrera D. Common biases in systems thinking. *Journal of Systems Thinking Preprints*. 2024. doi:10.54120/jost.pr000037.v1
32. Cabrera D. Hidden in plain science: Testing a universal ontology in the results of 13 scientific disciplines. *Journal of Systems Thinking*. 2026;6: 1–24.

33. Cabrera D, Cabrera L. Scale-free structure: Is knowledge just the recurring emergence of universal structures across scale? *Journal of Systems Thinking*. 2026;6: 1–40.
34. Cabrera, D., Cabrera, L., & Cabrera, E. A Literature Review of the Universal and Atomic Elements of Complex Cognition.
35. Cabrera D, Cabrera L, Cabrera E. Counterexamples and Resolutions Data. The DSRP counterexample record: candidate falsifications of O-Theory and the resolutions that settle them. 2026. Available: dsrpevidence.org/
36. Cabrera D, Cabrera L, Cabrera E. Convergent Independent Evidence Dataset. Cabrera Research Lab Living DSRP Evidence Compendium: independent convergent evidence for the universal and atomic elements of complex cognition [Data set]. 2026. Available: dsrpevidence.org/
37. Cabrera D. Knowledge Age Operating System: Four Principles of Organizational Design. Loveland, CO: Project N Press; 1996.
38. Cabrera D. A Mathematical Theory of Organization: DSRP as Universal Code of Mind and Nature for Organizing, Evolving, and Understanding Information. *Journal of Systems Thinking*. 2024;4: 1–23.

Appendix A. Secondary comparisons and replication

A1. Family-level pairwise comparisons

Table A1. Family-level D/S/R/P pairwise comparisons (exact McNemar with Holm correction). 'A only' and 'B only' count papers containing one Pattern of the pair but not the other; p values are exact McNemar tests with Holm correction across the six comparisons.

Pair	A only	B only	Raw p	Holm p	Result
D vs S	4	11	.118469	.236938	Not significant
D vs R	4	16	.011818	.035454	R > D
D vs P	11	0	.000977	.003906	D > P
S vs R	8	13	.383310	.383310	Not significant
S vs P	18	0	.000008	.000038	S > P
R vs P	24	1	.000002	.000009	R > P

A2. Strict-v10 pairwise comparisons

Table A2. Strict-v10 pairwise comparisons (exact McNemar with Holm correction). The same comparisons as Table A1 under strict DSRP-v10 coding, which requires both poles of a structure to be explicit in the analytical record.

Pair	A only	B only	Raw p	Holm p	Result
D vs S	4	7	.548828	1.000000	Not significant
D vs R	8	1	.039063	.117188	Not significant after Holm
D vs P	8	0	.007813	.031738	D > P
S vs R	11	1	.006348	.031738	S > R
S vs P	11	0	.000977	.005859	S > P
R vs P	1	0	1.000000	1.000000	Not significant

A3. Direct-role pairwise comparisons

Table A3. Direct LRL-role pairwise comparisons (exact McNemar with Holm correction). Each row compares the prevalence of two direct LRL roles across the 32 paired papers; no pairwise difference survives correction, consistent with the overall Cochran's Q result.

Pair	A only	B only	Raw p	Holm p
FB vs FT	6	3	.507813	1.000000
FB vs RV	12	6	.237885	1.000000
FB vs AL	8	5	.581055	1.000000
FT vs RV	10	7	.629059	1.000000
FT vs AL	5	5	1.000000	1.000000
RV vs AL	5	8	.581055	1.000000

A4. Data and replication

Two replication archives accompany this paper: one for the analytical-record analyses in Sections 3.1–3.3, one for the full-text replication in Section 3.4. Together they contain every codebook, instruction, and threshold used; the calibration set with its key and both candidate coders' answers; the complete verbatim harvest and every coder's exact input and output; and scripts that regenerate every number and figure

reported here. No coding decision or inferential test in this paper lies outside them. Source papers are excluded as copyrighted material; the manifest lists each one with the identifier needed to retrieve it.