Pattern-Recognition Unity (PRU): A Framework Specification

A Meta-Language for Grounding, Reference, and the Experience–Language Interface

Florin Cojocariu

2026-02-10

Contents

Abstract

This paper develops a meta-language for the relationship between experience and language, motivated by a structural diagnosis: natural language — including the precise technical language of philosophy — lacks stable constraints to hold apart the experiential and linguistic poles of cognition without collapsing one into the other. Every attempt to specify what ”grounding” or ”reference” is produces another linguistic description, which is exactly the kind of thing that needs grounding. When we ask what connects the word ”cat” to cats, every answer — causal chains, tracking relations, informational content — is itself a linguistic description of a connection living entirely *inside* language and *outside* the world, not the connection itself. The debate cycles because it uses the medium whose limitation it is about.

The meta-language is built within the Pattern-Recognition Unity (PRU) framework, which rests on two theoretical claims with empirical support. First, Pattern-Constellations: organisms form unified attractor states that integrate sensory input and motor-affective response as a single learned configuration, not as separate processes wired together. Second, Pattern-Recognition Unity: every recognition event is already a reaction — pattern and recognition are inseparable aspects of one event, not a detection step followed by a response step. These claims ground the notation in a substantive account of pre-linguistic cognition rather than in stipulation. The notational devices — marking the experiential pole without redescribing it in linguistic terms, marking the linguistic pole as linguistic, representing the coordination between them as a structured relation between different kinds of thing — are not arbitrary conventions but expressions of the architecture the framework describes. This is representational hygiene built on cognitive theory, not formalism for its own sake, though the framework is designed to support further formal development.

The paper proceeds in three parts. Part A diagnoses why a meta-language is needed. Part B develops the pre-linguistic foundation: Pattern-Constellations as unified attractor states integrating sensory input and motor-affective response as one event, the PRU thesis that pattern and recognition are inseparable aspects of a single event, composite constellations as the relational texture of experience, basic language as sound-patterns inside constellations, and the grounding of this architecture in Panksepp’s primary emotional systems. Part C introduces human language proper: the emergence of concept-words when labels operate in the absence of experience, the grounding operator that tethers linguistic patterns to experiential constellations, generics and predication as expressions and decompositions of unified patterns, and applications to classical problems of reference.

What This Is and How to Use It

This is not a paper in the classical sense but a foundational framework for cognition and language, built on two central ideas: Pattern-Constellations and Pattern-Recognition Unity. The strength of this framework becomes visible in its applications (for example, a critique of the Vector Grounding Problem or a mechanistic implementation of Radical Enactivism), but I kept running into a structural problem: it is effectively impossible to both introduce a new framework and apply it in a single paper without doing both tasks poorly.

This text is the solution. It serves as a reference point for future applications, where anyone interested can inspect the full scaffolding of PRU. I also hope it becomes a target for critique and improvement, because the theoretical ramifications and potential applications are wide enough that no single person can develop all of them.

To function as a reference text, the exposition is deliberately modular. Part B can be read independently as an account of basic minds and basic grounding. Part C, especially §§7–11, can be treated as the core reference for PRU's linguistic and grounding apparatus.

Paragraph numbering is continuous across parts, so that later papers can cite: §3 for Pattern-Constellations, §6 for x^o/x^c, §7 for LPCs, §8 for $\mathcal{R}$, §9 for generics, §11 for the LLM diagnostic, and so on.

Because the framework introduces a new meta-language, it also introduces new definitions. Each definition is boxed to make it easy to locate in the text, and all of them are collected at the end for quick reference.

PART A: THE PROBLEM

- **§1.1** diagnoses a structural difficulty in the grounding debate: every attempt to specify what grounding *is* produces a linguistic description of grounding, which is exactly the kind of thing the debate says is insufficient. This is the representational carousel.
- **§1.2** identifies the hidden assumption driving the carousel — that grounding is a property of representations — and reverses the direction: grounding runs from world to organism, not from language to world. The enactivist and ecological traditions that anticipated this reversal are acknowledged and situated.
- **§1.3** restates the problem as a question about two structurally different poles (experiential and linguistic) and the coordination between them.
- **§1.4** shows why the representational framework, with its template architecture, cannot serve as the solution, and motivates the need for a meta-language — not as an end in itself, but as the instrument that makes a substantive theoretical framework possible.
- **§2** goes through the most important related works.

1. The Carousel

1.1. Concepts About Grounding Are Not Grounding

How does the word "cat" connect to real cats? This question — the grounding problem in miniature — has organized three decades of debate in philosophy of mind, cognitive science, and AI. From (Harnad, 1990) symbol grounding problem through Searle's Chinese Room to the current dispute over LLMs[1], the question has the same architecture: on one side, a representation; on the other, the world; between them, a gap that theory must bridge.

The problem has force because representations, qua representations, relate only to other representations. A word connects to other words. A definition connects to other definitions. An internal state described as "carrying information about cats" connects to the description of cats, not to cats. This is the representational carousel: representations all the way down, the world never reached.

Mollo and Millière (2025) bring genuine progress to this question by distinguishing five notions of grounding — referential, sensorimotor, relational, communicative, epistemic — and showing that four of them fail to qualify as referential grounding because they never leave language. Sensorimotor grounding links one representation to another. Relational grounding is the carousel in its purest form: words linked to words. Communicative and epistemic grounding presuppose the content they are supposed to establish. Only referential grounding — the direct hook onto worldly entities — survives the elimination.

But consider what has actually happened. The four eliminated notions are not failed *kinds* of grounding. They are *concepts about grounding*: linguistic descriptions, formulated in publicly articulable prose, of how representations might relate to something. Each names a relation between representations. Each fails precisely because it is a description rather than the thing described. Concepts about grounding are not grounding — this is what the carousel argument shows.

Now ask: what is the status of the survivor? Referential grounding is supposed to be grounding itself — the real connection, not merely a concept about connection. But the moment it is specified — causal-informational relations, selection history, the function of tracking worldly features — it becomes another concept about grounding: another publicly articulable description, connected inferentially to other descriptions. The debate has produced a fifth entry for the carousel it taught us to recognize.

Mollo and Millière's own strongest candidate illustrates the difficulty. Their case for vector grounding rests on a structural isomorphism: geometric relations among word-vectors in an LLM's embedding space mirror relations among objects in the world. "Paris" is to "France" as "Berlin" is to "Germany" — and the vector space captures this. But the worldly relations that the isomorphism maps onto are themselves linguistically specified. "Capital of," "larger than," "part of" are descriptions of relations, not the relations themselves. The isomorphism thus holds between two representational structures: one internal to the model, one internal to our characterization of the world. This is relational grounding — the carousel in formal dress.

This is not a failure of execution. The problem is structural: natural language — including the precise technical language of philosophy — produces concepts about grounding, not grounding itself, because grounding, if it is real, is not a linguistic relation. Every attempt to specify *in natural language* what referential grounding *is* converts it into a *description* of what it is, and descriptions are exactly the kind of thing the carousel argument says is insufficient because descriptions are entirely inside natural language and cannot reach the world. To better understand this point, we can look to the first-order logic's "liar's

[1] A selected list of works on this sprawling subject:. . .

paradox" in a language L. Tarski showed that the solution is to define truth in a meta-language M, where the truth predicate is contained (Hodges, 2022). The problem of grounding is not similar in form but it's similar in *methodological predicament*: the target phenomenon isn't stably specifiable in the same register you're diagnosing.

The debate has cycled for thirty years because it keeps attempting to reach the world through the very medium whose limitation the debate is about.

1.2. Grounding Runs From World to Organism: Basic Grounding

Why does the carousel keep turning? The standard framing contains a hidden assumption: grounding is a property of representations — something language needs to acquire. The question "how does the word 'cat' connect to real cats?" places grounding in the linguistic domain. Its subject is the word. The word is what lacks grounding. The entire problematic runs from language outward: *symbol* $\rightarrow$ *???* $\rightarrow$ *world*.

But consider a cat tracking a mouse. The cat's mind connects to the mouse with precision: position, speed, trajectory are tracked; the body organizes around the encounter; salience and urgency radiate from the prey. No language is involved. No representations with truth-conditions are in play. Yet the connection between mind and world is as tight as any the grounding debate discusses.

This is not grounding in the standard sense — there is no linguistic content to ground. But it is a mind-world connection: a structured, stable, concern-laden engagement with a worldly entity, formed through the organism's history of encounter. Call it *basic grounding*: the pre-linguistic, pre-content connection between an organism and its world that the standard debate presupposes but never thematizes.

Basic grounding is what makes linguistic grounding possible. The child who learns the word "cat" does not use the word to reach cats. She already reaches cats — through basic grounding built from encounter-events. The word arrives and coordinates with a connection already in place. Grounding does not run from language to world. It runs from world to organism — basic grounding, the formation of experiential patterns through encounter — and then language coordinates with the result. The standard question inverts this order.

The representational framework reinforces the inversion: representations are items that bear content, content is evaluated for accuracy, accuracy is assessed in language. Everything the framework touches becomes a candidate for linguistic evaluation. It simply has no resources to describe the organism's pre-linguistic encounter with the world without converting it into a representation — just as language cannot describe pre-linguistic experience without converting it into language. The tool erases what it touches.

The insight that grounding runs from world to organism rather than from language to world has been articulated before, in different terms. Gibson's (1979) ecological psychology argued that organisms perceive affordances directly — structured possibilities for action in the environment — without constructing internal representations. The cat tracking the mouse is a Gibsonian scene: perception and action are coupled, and the "information" is in the organism-environment relation, not in a representation awaiting interpretation. Varela, Thompson, and Rosch (1991) generalized this into the enactivist thesis: cognition is not the manipulation of representations of a pre-given world but the enactment of a world through a history of structural coupling between organism and environment. Hutto and Myin (2013, 2017) pushed the point furthest, arguing that basic minds — minds without language — have no content at all: their engagement with the world is real, structured, and concern-laden, but it is not *about* anything in the sense that a sentence is about something. Content, on their account, arrives only with language and social practice. PRU inherits this diagnosis. What it adds is not the insight that grounding runs from world to organism — that is the shared enactivist commitment

— but a notational framework that makes the insight formally precise, allows it to be extended into natural language and, crucially, prevents it from collapsing back into representational language the moment it is theorized. The enactivist tradition has struggled with exactly this problem: its descriptions of contentless engagement are themselves contentful linguistic descriptions, and the representational framework recaptures them at the moment of articulation. The meta-language developed in this paper is designed to resist that recapture.

1.3. Two Poles and a Coordination

If grounding is not a property of representations but an event undergone by organisms — if the direction runs from world to organism rather than from language to world — then the reference grounding problem needs to be restated. What does it actually ask? Who are the mind actors once representations are gone?

(a) **The experiential pole**: the organism's stable pattern of encounter, constitutively subjective, individual, formed pre-linguistically through serial, locus-bound, concern-laden events. This is what basic grounding produces — what the cat has with the mouse, what the child has with cats before she learns the word. It does not cease when language arrives. Adult human minds continue to form and reshape experiential patterns through encounter, alongside and beneath their linguistic competence.

(b) **The linguistic pole**: the public, collective pattern of word-use, constitutively intersubjective, sustained through communal practice (Wittgensteinian language games inside a form of life). This is what natural language *is*. And this happens only in our minds.

(c) **The coordination**: the learned coupling between (a) and (b) — the process by which an individual's experiential pattern becomes coordinated with a community's word-use pattern[2]. This is what reference *is*.

These three are structurally different. (a) is subjective and cannot be transmitted through text. (b) is collective and *can* be transmitted through text. (c) is a relation between things of different kinds, not a link between two items within the same register.

The standard reference grounding problem collapses (a), (b), and (c) into a single question about representations. The result is that (a) — the constitutively subjective, phenomenological pole — becomes invisible. It is not denied; it is obscured by a framing that can only register it as one more representation among others.

1.4. Why a Meta-Language

What follows from this diagnosis is the need for a framework that can hold (a) and (b) apart — that can mark the experiential pole without converting it into a linguistic description, and mark the linguistic pole without letting it silently stand in for the whole. Natural language lacks stable terms to do this. How does one talk about pre-linguistic experience without converting it into linguistic description?

1.4.1. The problem with representations

The representational framework offers an apparent solution, but understanding why it fails is instructive. The framework inherits a template architecture from the computational theory of mind: cognition works by fitting inputs to stored specifications. A template is a detachable item — it exists independently of any instance, can

[2] At this point, this seems very close to what (Mollo and Millière, 2025) are doing with their isomorphism, but we'll see that the meta-language notation of PRU helps us understand where the connection with the real world happens.

be applied serially to different candidates, and divides them into fit and misfit. This structure automatically generates content-talk: the template becomes a content-object with satisfaction conditions, and recognition becomes the act of checking whether the world conforms to what is inside. “Aboutness” enters the story the moment a system is described as matching candidates to a specification.

But notice what the template architecture assumes: that pattern and recognition are two separate things — first the specification, then the check. Recognition is modeled as a serial process in which an independently given template is consulted and a verdict rendered. Part B of this paper will argue that this is not how recognition works. Pattern and recognition are one event: the organism settles into an attractor state in which sensing, motor readiness, and affective response fire as a unified configuration, not as a sequence of comparisons against a stored model.

This matters for the grounding problem because the template picture is what *generates* it. If cognition produces detachable content-objects, then each content-object needs connecting to a word, and each connection needs specifying — and we are back on the carousel. If cognition produces unified constellations in which the word-sound is already *inside* the pattern (as Part C will show), the specific gap that the grounding problem asks us to bridge does not open in the first place. The problem is not solved but preempted: it was an artifact of the template architecture, not a feature of mind.

1.4.2. Why a meta-language?

This does not mean the grounding problem disappears. Once language creates an autonomous realm — once concept-words operate outside the constellations they emerged from — a real coordination problem arises. But it is a different problem from the one the template picture generates, and it requires different tools. A meta-language is needed: a notation equipped with devices that indicate the experiential pole without redescribing it, mark the linguistic pole as linguistic, and represent the coordination between them as a structured relation between different kinds of things.

This is less than a formalism. A formalism derives results, generates theorems, performs calculations. The PRU framework, at this stage, does not do that. It is a meta-language: a language about the limitations of language, equipped with notational devices that mark where language reaches its edge.

It is also more than a metaphor. The distinction it draws has empirical traction and developmental predictions. It is not a suggestive analogy but a structural claim about the architecture of cognitive grounding.

But the meta-language is not the end point. It is what makes possible a substantive theoretical framework — built on Pattern-Constellations and Pattern-Recognition Unity — that characterizes pre-linguistic cognition, the emergence of language, and the coordination between them in terms that the representational framework systematically obscures. The notation serves the theory. The theory does the work.

Before the notation can be introduced, we need to describe what it is notation *for*. The experiential pole — what organisms have before and without language — must be characterized in its own terms. Part B does this: it develops the pre-linguistic architecture of minds that connect to the world through encounter, not through representation. Only then, in Part C, does language enter — and with it, the notation that makes the experiential-linguistic distinction precise.

A note on scope: this paper does not propose a complete metasemantics, a truth-conditional theory, or sufficient conditions for grounding. It is a target for empirical work that may lay the ground for a more formal approach, not a finished formal system.

2. Related Work

The PRU framework engages several active research traditions. This section does not attempt a comprehensive survey but positions the framework relative to the most relevant lines of work, specifying what PRU inherits, where it diverges, and what it claims to add.

2.1. Symbol Grounding

The grounding problem as a formal challenge originates with Harnad (1990), who showed that a purely symbolic system cannot ground its own symbols — that definitions in terms of other definitions produce an infinite regress. Searle's Chinese Room argument (Cole, 2026) provided the intuition pump: syntactic manipulation, however sophisticated, does not constitute semantic understanding. PRU inherits this diagnosis entirely. Where it parts company is with the framing. Both Harnad and Searle treat grounding as something symbols *lack* — a property that representations need to acquire through the right kind of causal or sensorimotor connection. PRU rejects this starting point. Grounding is not a property of representations. It is a structural feature of organisms that form experiential patterns through encounter. The question is not "how do symbols get grounded?" but "how does language coordinate with a grounding that is already in place?" This reframing — from symbol-outward to world-inward — reorganizes the problem rather than proposing a new answer to the old one.

2.2. Enactivism and Radical Enactive Cognition

The pre-linguistic architecture developed in Part B — basic minds, basic grounding, organisms that connect to the world through encounter rather than representation — occupies territory that the enactivist tradition has mapped extensively. (Varela et al., 2016) argued that cognition is not the manipulation of representations of a pre-given world but the enactment of a world through structural coupling between organism and environment. (Gibson, 2015) ecological psychology made a complementary case: organisms perceive affordances directly, without constructing internal representations. (Hutto and Myin, 2013) pushed the point furthest with Radical Enactive Cognition (REC), arguing that basic minds have no content at all — their engagement with the world is structured and concern-laden but not *about* anything in the propositional sense. Content, on their account, arrives only with language and social practice.

PRU is designed to be compatible with this tradition, not to compete with it. The enactivist diagnosis — that the representational framework distorts basic cognition by converting all engagement into content-bearing states — is one PRU shares entirely. What PRU adds is a notational and mechanistic articulation of distinctions that enactivists motivate at the level of phenomenology but do not formalize. Specifically: Pattern-Constellations provide a characterization of the "unit of skill" that REC leaves underspecified — what exactly is learned, stabilized, and redeployed in embodied coping? Attractor dynamics offer a mechanistic model for how these units form and operate. And the meta-language provides what the enactivist tradition has struggled to produce: a way of theorizing about contentless engagement without the representational framework recapturing the description at the moment of articulation. PRU aims to complete the enactivist picture, not to replace it.

2.3. Perceptual Symbol Systems

(Barsalou, 1999) perceptual symbol systems share PRU's central commitment to multimodal integration: cognitive representations are not amodal symbols but re-enactments of sensory, motor, and affective states.

PCs and perceptual symbols occupy similar theoretical space — both are distributed, multimodal, and grounded in bodily experience rather than in abstract symbolic structures.

The divergence is twofold. First, Barsalou retains the language of "representation" and "simulation": perceptual symbols are *simulations* of experience, activated during conceptual processing. PRU rejects simulation-talk because it preserves the template architecture in a subtler form — the simulation is still a detachable content-bearing state that can be "run" and evaluated. PCs are not simulations of past encounters but attractor basins shaped by them: the pattern is not replayed but re-entered, and the re-entry is already a reaction, not a depiction. Second, PRU makes the unity thesis central: pattern and recognition are one event, not a pattern (perceptual symbol) that is subsequently activated (simulated). This difference matters because it determines whether content-talk is generated by the architecture or not. Barsalou's perceptual symbols, being simulations with veridicality conditions, remain content-objects in the template sense. PCs are not.

2.4. Philosophy of Language

The classical problems of reference — (Kripke, 1990) causal chains, (Frege, 1892) sense-reference distinction, (Field, 1972) challenge to Tarski on primitive reference — are not targets PRU aims to defeat. They are frameworks PRU reinterprets using the meta-language developed in Part C. Kripke's baptism becomes the first operation of the grounding operator R on a new experiential-linguistic pair. Frege's cognitive difference between co-referential terms becomes a difference between linguistic pattern-constellations grounded in the same experiential attractor. Field's challenge — that Tarski's truth-definition presupposes but does not explain primitive reference — is addressed by showing that reference is not a primitive relation between word and world but the operation of R on co-developed experiential and linguistic structures. The detailed reinterpretations are developed in §7; what matters here is the stance: these are achievements to be preserved and reframed, not errors to be corrected.

2.5. Vector Grounding and the LLM Debate

Mollo and Millière (2025) are the immediate interlocutors. Their five-grounding tax^onomy brings genuine analytical clarity to the debate over LLM understanding, and their elimination of four candidates is — as §1 argues — essentially correct. PRU's contribution is upstream of their analysis: it provides a typed register framework that makes visible *why* four of the five groundings fail. They fail not because they are the wrong kind of relation, but because they have the wrong argument types for the grounding operator. Relational grounding stays inside ⟨x⟩-space. Communicative and epistemic grounding presuppose R rather than explaining it. Sensorimotor grounding was closest but needed a non-representational characterization of the experiential pole. The meta-language makes these diagnoses precise in a way that natural-language argumentation — including Mollo and Millière's own careful prose — cannot stably maintain, because the register distinctions collapse the moment they are stated without notation.

2.6. Affective Neuroscience

PRU grounds basic minds not in information-processing architecture but in Panksepp's (1998) primary emotional systems: SEEKING, FEAR, RAGE, CARE, PANIC/GRIEF, PLAY, and LUST. These subcortical systems, identified through electrical stimulation and lesion studies across mammalian species, function in PRU as modes of being that determine which Pattern-Constellations and composites are foregrounded at any moment. The choice is deliberate. Grounding the experiential pole in primary emotional systems rather than in sensory processing or computational architecture connects PRU to the growing "biological AI" current that argues intelligence requires homeostatic and affective stakes, not just information processing. It also makes a specific empirical commitment: the affective dimension of PCs is not an optional coloring added to a fundamentally cognitive structure but a constitutive feature of how the organism's world shows up. This commitment is developed in §4.3.

PART B: BASIC MINDS

Part B characterizes the experiential pole — what organisms have before and without language. §3 introduces Pattern-Constellations: unified attractor states integrating sensory, motor, and affective dimensions as a single learned configuration. §4 states the core thesis — Pattern-Recognition Unity — and argues that every recognition event is already a reaction, not a detection step followed by a response. §5 develops the pre-linguistic cognitive architecture: basic grounding as internal to each PC, composite constellations as the relational texture of experience, Panksepp's primary emotional systems as modes of being that configure which constellations are foregrounded, and basic language — sounds inside constellations that trigger pattern completion in conspecifics without reference, content, or grounding-as-a-relation.

3. Pattern-Constellations

3.1. Definition

Definition: Pattern-Constellation (PC)

A **Pattern-Constellation** (PC) is a unified attractor-like state in neural space integrating multiple pattern types: sensory patterns (visual, auditory, tactile, olfactory, gustatory), motor patterns (actions afforded and executed), affective patterns (emotional responses), social patterns (how others react, norms, joint attention), and interoceptive patterns (hunger, fatigue, arousal).

$$\left\{\begin{array}{l}\text{sensory patterns}\\ \text{motor patterns}\\ \text{affective patterns}\\ \text{social patterns}\\ \text{interoceptive patterns}\end{array}\right\}$$

Notation: $\{A\}$ where A is capitalized.
Reads as: "the pattern-constellation A."

$\{A\}$ is not a local object in the brain and not a simple feature map. It is an analytical notion: a way of talking about a stable energy configuration of the entire network, distributed over many units. We never encounter $\{BIRD\}$ or $\{MOUSE\}$ as such. We encounter *this bird here*, *that mouse there*.

In physiological terms, $\{A\}$ can be seen as a set of different regions in the brain that may activate together as a result of a specific sensing pattern, involving visual, auditory, motor, and feeling regions in the brain.

But if we say "the brain matches some actual experience against the template $\{BIRD\}$," we mislead ourselves. There is no separate template sitting on a shelf. All learned PCs are juxtaposed in the network as different ways it can resonate. When a fragment of a situation has the right structure, the network does not look up $\{BIRD\}$. It falls into a $\{BIRD\}$-like state as a whole:

$$\partial\{BIRD\} \rightarrow \{BIRD\}_{active}$$

This is one dynamical event, not a serial comparison.

Definition: Pattern Recognition Event

A pattern-recognition event $E(\{A\})$ is the activation of a pattern constellation $\{A\}$ as a result of certain internal or external stimulus that drives the entire neural system into the learned attractor. When we measure neurological activity in specific situations, we see the activation events which are the only empirical acces we have to pattern-constellations.

PCs are not templates; they are whole-network energy states. Within a resonant state, a portion of the current experiential field is treated as a whole *instance* — "this bird here," "that mouse there." The instance is what emerges when the PC is active, not a separate entity that the PC is matched against.

3.2. Neural Implementation

The attractor dynamics studied in Hopfield-type networks (Hopfield, 1982) with Hebbian learning provide a natural model for PCs. During learning, extensional encounters (*this* mouse, *that* mouse, another mouse) incrementally reshape synaptic weights, gradually sculpting an energy landscape in which stable basins correspond to learned constellations. During recognition, mouse-like inputs drive the system into the appropriate basin, yielding content-addressable memory and pattern completion. (Hopfield, 1982)

The attractor dynamics support completion from partial cues, tolerance to noise, and graceful degradation. Unlike simple threshold devices, the basin's topology embodies a standing capacity to pull diverse inputs toward unified type-level responses. Software models displaying this few-shots pattern-learning ability are available.

3.3. Pattern-Constellations as Intensionality-Constructors

Philosophically, PCs provide a mechanism by which intensionality emerges at a pre-verbal level. The classical intension–extension distinction seems to have a prior realization in pattern recognition dynamics.

The intensional work is not done online at recognition time. It is embodied in the pre-sculpted landscape of the network. Before learning, an organism's experience is purely extensional: a succession of untyped instances. After a sufficient number of exposures, learning occurs, and the attractor structure enables type-level responses: many different inputs are treated as the same kind of thing. The basin's topology just *is* a standing intensional capacity.

This reframes intensionality as a property of pattern dynamics that language later exploits, rather than as a purely linguistic or symbolic phenomenon. We may call this "basic intensionality", in line with a practice in this paper to prefix "basic" all cognitive concepts that are pre-linguistic.

Conversely, what we call "recognition" in PRU is a unified event of activating a PC when experiencing something that this PC resonates with. It is an instantiation: a move from the intensionality encapsulated in the PC to the experienced extensionality of an instance. We never "see" the PC or the low-energy basin. We only ever meet *this mouse*, *that bird*, *that person*, as experienced PRU events, while the intensional structure hums underneath as a low-energy attractor.

3.4. Properties of PCs

- **Integration**: All components of $\{A\}$ tend to co-activate — sensory, motor, affective, interoceptive dimensions fire as one complex.
- **Partial activation**: A fragment $\partial\{A\}$ can drive full $\{A\}$ activation (pattern completion).
- **Individual variation**: $\{A\}$ varies between individuals based on learning history, environment, frequency, and type of encounters. There is no "objective" $\{DOG\}$ hovering above individual experience. $\{X\}$ exists only as someone's $\{X\}$.
- **Overlap**: $\{A\}$ and $\{B\}$ can partially share structure, supporting analogy and category relations.
- **Resonance**: PCs are resonators, not templates. The system "falls into" an $\{A\}$-state rather than comparing inputs against stored models.

4. Pattern-Recognition Unity

4.1. The Core Thesis

A pattern-recognition event $E(\{A\})$ is already a *reaction*:

- sensory fields are organized in an A-like way,
- motor dispositions are shaped accordingly (approach, avoidance, grasping),
- affective responses are set (fear, curiosity, joy),
- interoceptive state is modulated.

All these elements are learned together during training. The constellation $\{A\}$ is not only a sensory cluster but a reaction pattern too: a motor–feeling–interoceptive complex that tends to arise as one unit. When an $E(\{A\})$ episode occurs, the organism does not first compute "this matches A" and then decide what to do. The "match" and the "doing something about it" are the same event. Motor centers activate at the sight of prey or the sound of a predator even when no overt action follows — the full reaction pattern fires, with downstream inhibition selecting which components are expressed.

This is the core of:

> **Definition: Pattern Recognition Unity**
>
> **Pattern Recognition Unity** (PRU) principle states that every pattern-recognition event is a unified pattern–reaction event. "Recognizing A" is the system's way of *being ready to react to-A*. There is no "reaction process" separate from the "recognition process".

This explains why recognition is immediate rather than serial. PCs act as resonators, not templates. The network does not iterate over stored patterns and compare them to the input. Instead, a partial presence $\partial\{A\}$ in the current state can drive the whole constellation into activation, bringing with it the associated reactions.

Notation: $E(\{A\})$ marks an experiential pattern-recognition event — the network falling into an $\{A\}$-like attractor state. E stands for experiential actualization. These events are:

- **Atomic**: cannot be decomposed into separate "pattern" and "recognition" steps.
- **Subjective**: occur in individual experience.
- **Pre-linguistic**: present in animals and pre-linguistic infants.
- **Reactive**: perception and response are aspects of one event.

4.2. Pattern-Matching vs. Pattern-Recognition

A distinction is essential here, because the technical literature (and by extension, the philosophical one) uses "pattern recognition" to name something fundamentally different from what PRU means by it.

Pattern-matching is an algorithmic process: an input is compared against stored templates or feature vectors, and the best match is selected. It is serial through logical dependency — each step cannot begin until the previous step completes. Input arrives, is preprocessed, is compared against a library of candidates, and a label is returned. The “pattern” and the “recognizer” are independent entities that meet in computation. It is trying keys in a bunch until one fits the lock.

Pattern-recognition in PRU’s sense is a resonance event. There is no independent input, no independent template, no comparison step. The system falls into an attractor state as a whole. The “pattern” and the “recognition” are not two things that meet — they are two aspects of one event. There is no moment at which the pattern exists without being recognized, and no moment at which recognition occurs without a pattern to organize it. This is recognizing an entire tune just by a few notes, or a face in a crowd just by a couple of traits.

Pattern-matching is very useful for building computational systems (algorithms and reasoning seem entirely based on them) but should not be confused with PRU’s notion of pattern-recognition as a single, resonant pattern–reaction event. The confusion between the two is endemic in the cognitive science literature and produces systematic distortions when computational pattern-matching is used as a model for biological pattern-recognition. In particular, it generates the illusion that recognition is a two-stage process (detect, then react) when it is in fact a unified event.

Pattern matching and representational theories, the template matching process as the implicit assumption in representational theories.

5. The Structure of Basic Minds

5.1. Basic Grounding Is Internal to PCs

What we called “basic grounding” in §1.2 — the pre-linguistic, pre-content connection between an organism and its world — does not require a relation between the organism and something external. It is internal to each pattern-constellation.

Because {MOUSE} integrates sensory input and motor-affective response as a single learned configuration, the cat’s connection to mice is constituted by the constellation itself. The cat does not first detect the mouse and then connect to it. Detection *is* connection activation: the network’s falling into a {MOUSE}-state is simultaneously the organization of sensory fields around the mouse, the activation of stalking-readiness, and the SEEKING-ex^citement. The grounding is the constellation.

We can write this analytically as:

$$\{MOUSE\} = \{MOUSE_{input}, MOUSE_{response}\}$$

— but this decomposition is an analytical convenience, not a description of two things wired together. The input and response dimensions of {MOUSE} were learned together, they activate together, and they constitute a single attractor basin. The grounding is not a relation added to the PC; it is what the PC *is*.

This is why basic grounding requires no “self” entity and no explicit coordination relation. Each PC already contains the organism’s way-of-being-toward the thing it is a constellation of. {MOUSE} already contains the cat-ness — the specific motor-affective profile that makes this a *cat’s* mouse-constellation rather than an owl’s or a snake’s.

5.2. Composite Constellations: The Relational Texture of Experience

A cat tracking a mouse does not merely activate {MOUSE} in isolation. The mouse appears *in the grass*, near a *shadow* overhead, at a particular *distance from shelter*. The relational texture of the cat's experienced world is captured not by relations between independent PCs, but by composite constellations:

$$\{\{MOUSE\}, \{GRASS\}\}$$

This composite PC has its own integrated response profile — the specific motor-affective pattern of mouse-hunting-in-grass, which may include but may differ from mouse-hunting-in-open-ground or mouse-at-a-distance. The grass doesn't just co-occur with the mouse; the *composite* {MOUSE-in-GRASS} is a learned configuration with its own attractor dynamics, and may have (or not) its own motor readiness (low crouch, slow approach through cover), and its own affective character (like a danger feeling associated with it).

Composition is how the animal's world acquires relational structure:

- $\{\{MOUSE\}, \{GRASS\}\}$ — prey in cover, with cover-specific stalking responses.
- $\{\{SHADOW\}, \{OVERHEAD\}\}$ — aerial threat, with freeze/flatten responses.
- $\{\{CRY\}, \{OFFSPRING\}\}$ — distress from young, with approach/nurturing responses.

Each composite is a unified PC at a higher level. It is not a proposition ("the mouse is in the grass") and not a belief. It is a learned configuration that organizes the organism's engagement with a *situation*, not just an object. The animal's experienced world is a dynamic, ever-changing field of such composites, layered and overlapping, each carrying its own response profile.

5.3. Modes of Being: Panksepp's Primary Emotional Systems

Which composites are salient at any moment depend on the organism's mode of being. A hungry cat and a satiated cat encounter the same scene differently — not because they see different things, but because different constellations and composites are foregrounded.

Panksepp (1998, 2011) identified seven primary emotional systems through electrical brain stimulation, pharmacological challenge, and lesion studies across mammalian species (Davis and Montag, 2019):

System	Valence	Function
SEEKING	+	Exploration, foraging, anticipation
FEAR	-	Avoidance of danger
RAGE	-	Response to restraint/frustration
CARE	+	Nurturing behavior
PANIC/GRIEF	-	Separation distress
PLAY	+	Social joy
LUST	+	Sexual pursuit

Key empirical findings: these systems are subcortical and conserved across mammals — a decorticate rat still shows FEAR, RAGE, SEEKING behaviors (Montag and Davis, 2018). Electrical stimulation directly evokes behavioral-emotional patterns. Animals will work to activate positive systems and work to deactivate negative ones. The systems interact: SEEKING energizes pursuit across all domains; FEAR and RAGE are adjacent and can flip; PANIC inhibits PLAY.

These are not "drives" added onto perception. They are *modes of being* that determine how the world shows up. Each mode reconfigures which PCs and composites are foregrounded:

- In SEEKING mode, $\{\{MOUSE\}, \{GRASS\}\}$ is salient — the environment is experienced as a field of potential prey.
- In FEAR mode, $\{\{SHADOW\}, \{OVERHEAD\}\}$ dominates — the same environment is experienced as a threat-landscape.
- In CARE mode, $\{\{CRY\}, \{OFFSPRING\}\}$ foregrounds — the environment is a field of nurturing affordances.

The same scene, the same PCs, but different modes of engagement foreground different composites with different response profiles. This is not a change in what is perceived. It is a change in *how things matter*.

The organism's interoceptive state (hunger, fatigue, arousal, hormonal condition) determines which mode is foregrounded. Interoception is not input *to* the organism — it *is* the organism's current mode of being. Hunger is not information about a need. It is the need itself, experienced as a reconfiguration of what matters.

Modes can be modeled as specific subspaces inside the complex, multidimensional $\sum X_i$ space.

5.4. Basic Language: Sounds Inside Constellations

Many animals produce characteristic sounds as part of their pattern-recognition events. A chicken seeing a hawk produces a specific alarm call. A vervet monkey seeing a snake produces a different call than when seeing an eagle (Seyfarth et al., 1980). These vocalizations are motor-sensory patterns — larynx, breath, auditory self-monitoring — that are *inside* the relevant PC, not referential labels attached to it.

The chicken's hawk-alarm is part of {HAWK} in exactly the same way that freezing or looking up is part of {HAWK}. It is a motor component of the pattern-reaction event. The call does not *refer to* hawks. It does not *mean* hawk. It is part of what the hawk-constellation *is* for this organism — one dimension of the unified pattern-reaction alongside the visual vigilance, the freeze response, and the FEAR affect.

When another chicken hears the alarm call, the sound functions as a partial cue: $\partial\{HAWK\} \rightarrow \{HAWK\}_{active}$. The hearing triggers pattern completion — the full hawk-constellation activates, including the visual scanning, the motor readiness, and the affective state. This is not decoding a message. It is resonance: a fragment of the pattern (the auditory component) drives the whole constellation into activation.

This is *basic language*: communication without reference, without content, without grounding-as-a-relation. The sound is inside the pattern, and pattern completion does the rest. The result is mimetic coordination in a group, one that was selected to increase survival chances for the species.

5.5. Self as Derived Structure

The totality of response dimensions across all of an organism's PCs constitutes what we might call the organism's identity. Every {X} contains motor-affective responses — the organism's way of being toward X. If we project every PC in the organism's repertoire onto its response dimensions, we get a structure that is invariant across changing objects of encounter: what remains constant as {MOUSE}, {SHADOW},

$\{$WARM-SPOT$\}$ come and go. This is what "self" means at the basic-mind level — not a separate PC sitting alongside others, but the common response-dimension running through all constellations. (Thinking at it, it is obvious that cats "characters" differ by their different reaction modes to the same stimulus.)

This structure is real and important, but developing a full account of it is beyond the scope of this paper. What matters here is that basic grounding does not require a self-entity. Each PC is already grounded by its own internal structure — the integration of input and response; basic grounding is a fundamental *internal* feature of pattern constellations. The self is what emerges when you look at the *totality* of this grounding across the organism's whole repertoire. (It must be underlined that we do not posit here that this *basic self* is limited to the sum of responses; something more may act on the totality of responses to make them biased one way or another, but this is beyond the scope of this paper. The point is that we can model *the self* as distributed in pattern constellations and not a ghost in the machine.)

5.6. Summary: What Basic Minds Have

An organism with basic grounding has:

1. **Stable PCs** $\{X\}$ formed through encounter, each integrating sensory input and motor-affective response as a single configuration.
2. **Composite PCs** $\{\{X\}, \{Y\}\}$ that capture the relational structure of the experienced world.
3. **Panksepp's primary emotional systems** as modes of being that foreground different composites.
4. **Basic language**: sounds inside PCs that trigger pattern completion in conspecifics.
5. **A derived self-structure**: the totality of response dimensions across all PCs.

What it does not have: any separation between experience and language. Everything is inside PCs. There is no autonomous linguistic realm, no concept-words, and therefore no grounding *problem*. The problem arises only when language creates an outside — which is where Part C begins.

PART C: LANGUAGE AND GROUNDING

Part C introduces human language and the notation that makes the experiential-linguistic distinction precise. §6 traces the developmental trajectory from basic language to human language: the integration of word-sounds into experiential constellations (object-words, x^o), the emergence of concept-words (x^c) when labels operate in the absence of experience, and the creation of an autonomous linguistic realm. §7 defines Linguistic Pattern-Constellations ($\langle x \rangle$) — patterns of word-use in linguistic space — and establishes the constitutive asymmetry between experiential and linguistic poles. §8 introduces the grounding operator $\mathcal{R}$, characterizes its formal properties, demonstrates what grounding looks like as a living, encounter-driven process, and applies the operator to reframe classical problems of reference (Kripke, Frege, Field). §9 develops unified patterns and generics as cognitively primitive structures from which predication emerges by decomposition, inverting the standard assumption that predication is the foundation and generics a defective form of it.

6. From Basic Language to Human Language

6.1. The Continuity

Human language begins the same way as basic language. The child's first words are motor-sensory patterns — movements of mouth, tongue, and breath, coupled with auditory self-monitoring — that become integrated into existing PCs through encounter. When the child repeatedly hears "dog" while experiencing dogs, the sound-pattern enters $\{DOG\}$ as a motor-sensory component, alongside the visual, tactile, and affective dimensions already present. This picture is consistent with the usage-based account of language acquisition (Tomasello, 2009), in which children's earliest words function as parts of holistic action-routines rather than as referential labels attached to pre-existing concepts.

At this stage, the word functions exactly like the chicken's alarm call. "Dog" said while seeing a dog is part of $\{DOG\}$ — one more dimension of the pattern-reaction event. It is not yet a label *for* the dog. It is part of what experiencing-a-dog *is* for this child. Call this the **object-word** function: x^o.

Definition: Object-Word

Notation: x^o — the object-word. A motor-sensory pattern (the pronounced and heard word-form) functioning as a component of a PC and as an ostensive handle for it. It is employed in natural language in ostension, in the direct presence of the object that it labels.

$\{DOG, dog^o\}$ marks a constellation in which the word-sound "dog" has been integrated. In $E(\{DOG, dog^o\})$ events, saying or hearing "dog" is part of the dog-encountering event. dog^o labels *this dog here* — the current activation of $\{DOG\}$.

6.2. The Exit: Concept-Words

Something happens in human development that has no parallel in animal communication. The child begins hearing "dog" in contexts where no dog is present — in stories, questions, corrections, classifications, games. "Where's the dog?" "Dogs are animals." "Is that a dog or a wolf?"

In these contexts, the word operates without any current $E(\{DOG\})$. It participates in patterns of language-use that have their own regularities — distributional, syntactic, inferential, pragmatic. Over time, these regularities become their own kind of pattern: a pattern not of dog-encounters but of dog-talk.

This is where the **concept-word** emerges: x^c.

Definition: Concept-Word

Notation: x^c — the concept-word. A word functioning within linguistic space, labeling patterns of use in sentences rather than experienced instances; the concept-word dog^c can be seen as labeling the totality of sentences containing "dog", known to a user [a]. As a token it is indistinguishable from the object-word; only context makes the difference.

[a] There is still s subjective side to it, explaining why concepts may be understood differently from user to user.

x^c does not have "experiential instances" in the way x^o does. It operates over how "x" is used across utterances, constructions, inferences, and language-games. When the child can answer "What is a dog?" with "An animal that barks," she is using dog^c — operating in linguistic space, not labeling an experiential instance. Its instance is always *a sentence*.

The developmental sequence:

$$\{DOG\} \rightarrow \{DOG, dog^o\} \rightarrow \{DOG, dog^o, dog^c\}$$

Stage 1 is Part B's material: a pure experiential constellation. Stage 2 integrates the word-sound as part of the constellation — basic language. Stage 3 is the uniquely human achievement: the same word now also operates *outside* $\{DOG\}$, in an autonomous linguistic space. There is evidence that infants initially use words to label experienced instances before deploying them in the absence of referents — a transition broadly consistent with the shift from x^o to x^c use described here (Bloom, 2001), (Tomasello, 2009).

The object/concept aparent dichotomy has a long history and probably the most striking can be found in Carnap's Aufbau: $>$ It makes no logical difference whether a given sign denotes the concept or the object, or whether a sentence holds for objects or concepts. There is at most a psychological difference, namely, a difference in mental imagery. Actually, we have here not two conceptions, but only two different interpretative modes of speech. Thus, in construction theory we sometimes speak of constructed objects, sometimes of constructed concepts, without differentiating. (Carnap, 1969, p. 10)

The notation above is just formalizing an idea that was floating around for a long time.

6.3. Language's Power and the Grounding Problem

Something that was *inside* $\{X\}$ — the word-sound, functioning as a motor-sensory component — now also operates *outside* $\{X\}$, in a realm of linguistic patterns. This is the decisive moment. A component of the experiential constellation has acquired an autonomous life. It can be used, combined, negated, questioned, embedded in conditionals — all without any current experiential activation.

This is both language's power and the source of the grounding problem. The power: you can talk about justice, unicorns, the past — things that have no current $E(\{X\})$ to activate. The source of the problem: the autonomous linguistic realm can drift free of experience entirely.

Definition: Linguistic mediation

Notation:: $\exists(\{A, a^c\})$ marks operations where the PC is being talked about or manipulated at the conceptual level without any current instance present. "Dogs are animals" uses dog^c: $\exists(\{DOG, dog^c\})$. "There is snow on the mountain" said while looking at a map uses $snow^c$ without a current $E(\{SNOW\})$. $\exists$ marks the appearance of language as an autonomous realm.

6.4. A Worked Example: 'Dog

The experiential pole.

Language enters.

The child hears "dog" while experiencing dogs. The sound-pattern integrates into the constellation: $\{DOG, dog^o\}$. Now experience has a word labelling it: $E(\{DOG, dog^o\})$. This is an experiential event in which the word-sound is *inside* the constellation, functioning as one more sensory-motor component alongside the fur and the barking. dog^o labels *this dog here* — it is an instantiator tied to the current activation of $\{DOG\}$.

Then the child starts hearing "dog" without dogs present. "Where's the dog?" "Dogs are animals." The word starts operating in linguistic space — participating in patterns of use that have their own regularities. dog^c emerges: the concept-word, a node in a network of inferences, constructions, and language-games. The same phonological token now functions in two registers. dog^o labels instances. dog^c labels the linguistic pattern.

See what the notation prevents here: it prevents you from treating "dog" as a single item that somehow does both jobs at once. The superscripts force you to specify which register the word is operating in at any given moment. Without them, the sentence "the word 'dog' refers to dogs" looks like it says one thing. With them, you see it conflates two operations: dog^c operating in linguistic space, and dog^o anchoring to an experienced instance.

The full architecture.

Now the child is a competent speaker. She has $\{DOG\}$ — the experiential attractor basin. She has $<dog>$ — the linguistic pattern-constellation, the distributional-inferential profile of "dog" across her language use. These are different kinds of structure: $\{DOG\}$ is subjective, individual, non-transmissible through text; $<dog>$ is collective, approximately shared, transmissible through text.

$\mathcal{R}(\{DOG\}, <dog>)$ marks the grounding operator tethering the linguistic pattern to the experiential constellation. When the speaker says "dogs are loyal" in a way that *means something to her* — not merely repeating a sentence she's heard but speaking from her experience of dogs — R is operating. The concept-word dog^c activates $<dog>$, and R co-activates $\{DOG\}$, and the word is grounded.

But what $\exists(<dog>)$ *without* $\mathcal{R}$, is? This is a system that navigates the linguistic pattern — knows that "dog" co-occurs with "loyal," "pet," "bark," "fetch" — without any experiential attractor basin to tether it to. This is what an LLM has. It is not missing a component that could be bolted on. It is missing the developmental history that would have made $\mathcal{R}$ possible: the stage at which dog^o formed *inside* $\{DOG\}$, from which dog^c could later emerge and acquire an autonomous life while remaining coupled to its experiential origin.

The notation does not deliver $\{DOG\}$ to the reader. It marks where $\{DOG\}$ would be, so that its absence — in a system, in an argument, in a philosophical theory — becomes visible as lack of grounding rather than silently elided.

6.5. A Note on Acquisition

The staged picture is an analytical simplification. In real development, the child is immersed in adult language use from the beginning. Consider "apple." When the child learns "apple," it is presented with the object and, over repetitions, forms $\{APPLE\}$. This happens not only when adults repeat "apple" and point, but also when the child observes adults interacting with apples and talking about them. In Wittgenstein's terms, the child is already immersed in adult practices and language games, even without understanding them.

For this reason, it is impossible to specify exactly when a^c is acquired, and it is misleading to treat the acquisition as a clean step. The transition to x^c is not a discrete event but a gradual process that begins as the child participates in more practices and games, the label-like role slowly lifting into the concept role.

In the analytical model, however, the distinction between x^o and x^c remains useful. There is evidence that infants use words as labels for things before they use them as fully developed concepts (Bloom, 2000; Tomasello, 2003). By the time x^c is operational, x^o still carries its original connection to experience and now functions as an instantiator: a way of binding the intensional structure of $\{X\}$ to a concrete instance, usually marked by articulation ("the x," "this x," "an x").

7. Linguistic Pattern-Constellations (LPCs)

7.1. Motivation

Language in use builds its own pattern-constellations. When we say that for a speaker "the word 'justice' has a certain feel and inferential profile," we are pointing at a *linguistic* pattern: a pattern over sentences, arguments, narratives, and dialogues. These patterns live in the abstract $\exists$-realm rather than directly in the experiential E realm.

7.2. Definition and Notation

Definition: Linguistic Pattern-Constellation

Notation: $\langle a \rangle$ — angle brackets. An LPC $\langle a \rangle$ is the pattern of how the word "a" behaves in linguistic space — where it appears, in which constructions, with which neighbors, in which moves (assertion, question, imperative), in which genres. It is built from many $\exists(\{A, a^c\})$ events and from co-occurring words, syntactic frames, discourse roles, and pragmatic functions. **Reads as**: "the Linguistic Pattern-Constellation for the word 'a'."

x^c is the concept-word that labels $\langle x \rangle$. The relation between x^c and $\langle x \rangle$ parallels the relation between x^o and $\{X\}$: just as x^o labels instances of the experiential constellation, x^c labels the linguistic pattern and appears in its instances, i.e. in sentences.

x^c can be better understand in the Wittgensteinian frame as labeling the totality of uses of an word. This view also stresses out that concept-words are learned in language-games connected to certain forms of life (that is common practices), and not by employing some sort of dictionary.

7.3. Hierarchy of Patterns

- **Level 1: Primitive PCs** — $\{A\}$ or $\{A, a^o, a^c\}$ — constellations integrating sensory, motor, affective, and linguistic components.
- **Level 2: LPCs** — $\langle a \rangle$ — patterns of word-use in linguistic space.
- **Level 3: Higher structures** — narratives, theories, ideologies, scientific frameworks, built from LPCs and their interrelations.

7.4. Experience and Existence: E vs $\exists$

Primitive PCs can have $E(\{A\})$ and $E(\{A, a^o\})$ — direct experience — and participate in $\exists(\{A, a^c\})$ once a^c appears.

LPCs have no direct $E(\langle a \rangle)$, because there is no direct sensory experience of "the pattern of use of 'a' " as such. They exist purely in the $\exists$-realm: $\exists(\langle a \rangle)$. But they are ultimately grounded through chains of $E(\{A, a^o\})$ episodes that gave rise to the words and constructions in which "a" appears.

Primitive PCs link language to the world and action. LPCs describe the structure of language itself. The two levels are coupled by a^o and a^c, and by the fact that many LPCs include words whose x^o uses are grounded in $E(\{X\})$.

7.5. Subjective vs Collective: The Constitutive Asymmetry

A point the notation makes visible but that natural language habitually obscures:

$\{X\}$ is *constitutively subjective*. It is formed through the individual organism's history of encounter — serial, locus-bound, laden with affect and concern. My $\{DOG\}$ and your $\{DOG\}$ differ because our encounters with dogs differ. $\{X\}$ exists only as someone's $\{X\}$.

$\langle x \rangle$ is *constitutively collective*. It is the pattern of use of the word "x" across a community of speakers. It is publicly articulable, transmissible through text, reconstructable from a corpus. $\langle dog \rangle$ is approximately the same for all competent English speakers — not because their experiences are the same, but because the linguistic pattern is a social object, maintained by use.

This asymmetry is what makes communication both possible and imperfect. Possible, because $\langle x \rangle$ provides a shared coordinate system. Imperfect, because $\{X\}$ is always individual, and no amount of $\langle x \rangle$-level precision can fully transmit one person's experiential constellation to another.

8. The Grounding Operator $\mathcal{R}$

With both poles now defined — experiential PCs $\{X\}$ and linguistic PCs $\langle x \rangle$ — we can characterize the operator that holds them together. (The exposition here follows the logic of the framework rather than strict developmental order. Developmentally, the grounding operator is at work before LPCs are fully formed — but its characterization requires both poles to be in place.)

8.1. Why Grounding Requires an Operator

In basic minds, there is no grounding problem because there is no separation between experience and language. The chicken's alarm call is *inside* {HAWK}. There is nothing to connect, nothing to ground, nothing to bridge.

In human cognition, x^c has exited $\{X\}$ and operates autonomously. The patterns of word-use form their own structures — the LPCs defined in §6. These structures live in a different realm from experiential PCs: they are collective where PCs are subjective, publicly articulable where PCs are first-personal, transmissible through text where PCs are locus-bound.

The grounding operator $\mathcal{R}$ is what keeps the autonomous linguistic realm tethered to experience. It operates on what was inside $\{X\}$ and what is now outside it.

$$\mathcal{R}(\{X\}, \langle x \rangle)$$

This is not a mysterious bridge between word and world. It is the residual coupling between two structures that *co-developed* through learning: the experiential constellation $\{X\}$ and the linguistic pattern $\langle x \rangle$ grew together as the child's x^o was gradually lifted into x^c-use. They were never separate things that needed connecting. They are two aspects of the same developmental history that have become partially autonomous.

8.2. What $\mathcal{R}$ Is

$\mathcal{R}$ is an operator, not merely a relation that holds between its arguments. A relation describes a state of affairs — "$\{X\}$ and $\langle x \rangle$ are coupled." An operator *produces* something when applied to its arguments. $\mathcal{R}$ produces grounding: it takes an experiential constellation and a linguistic pattern and yields a grounded word — a word tethered to experience.

More precisely: $\mathcal{R}$ is what the cognitive system *does* when it activates $\{X\}$ and $\langle x \rangle$ together. The co-activation is not passive co-occurrence. It is an operation that resolves a word's potential for x^o and x^c use into a specific grounded configuration. Before $\mathcal{R}$ operates, x^c floats in linguistic space and x^o labels an instance. When $\mathcal{R}$ operates, the concept-word is tethered to the experiential constellation: the speaker *means something* by it.

This is why "operator" and not "relation": $\mathcal{R}$ is something the cognitive system performs, not something it has. It performs it each time a word is used with grounding — each time x^c activates $\{X\}$ or x^o invokes the conceptual space of $\langle x \rangle$.

In formal terms, $\mathcal{R}$ is an operator more in the way energy (the Hamiltonian $\hat{H}$) is an operator in QM rather than a classical operator in mathematics, in the sense that it renders a measure as an eigenvalue in a specific application. This has intriguing applications in modelling how meaning "collapses" in specific choices of words in a sentence, but it ex^ceeds the scope of this paper.

What $\mathcal{R}$ is *not*:

- **Not causal**: $\{X\}$ does not *cause* $\langle x \rangle$ or vice versa. They co-form through the same learning history.
- **Not a bridge between pre-existing separates**: $\{X\}$ and $\langle x \rangle$ co-developed. $\mathcal{R}$ is the operation that maintains their coupling as they become partially autonomous.
- **Not a matching operation**: $\mathcal{R}$ does not compare $\{X\}$ against $\langle x \rangle$. It *activates* their coupling — resolving a word's potential for x^o and x^c use into a specific grounded configuration.

The operation is enabled by a structural coupling that was established through Hebbian co-occurrence learning during $E(\{X\}, x^o)$ events. But the coupling is the *condition* for the operation, not the operation itself.

8.2.1. Formal Characterization

For a cognitive system S, $\mathcal{R}(\{X\}, \langle x \rangle)$ can operate iff:

1. $\{X\}$ is a stable attractor basin in S's experiential state space.
2. $\langle x \rangle$ is a stable pattern of linguistic/conceptual activation in S.

3. Activation in the $\langle x \rangle$ region reliably co-activates (or makes accessible) the $\{X\}$ basin, and vice versa.
4. This coupling was established through Hebbian co-occurrence learning during $E(\{X\}, x^o)$ events.

8.3. The Fundamental Identity: R as Grounding Operator

Our notation using object-words, concept-words and concept-properties is a shortcut for how $\mathcal{R}$ operates on pattern-constellations (including LPCs):

$$\mathcal{R}(x^c, x^o) \equiv \mathcal{R}(\{X\}, \langle x \rangle)$$

Level	Notation	What it captures
Surface (word-level)	$\mathcal{R}(x^c, x^o)$	Apparent “reference” in language use
Deep (constellation-level)	$\mathcal{R}(\{X\}, \langle x \rangle)$	Actual architectural coordination

This identity is not just notational equivalence — it is the structural explanation for the phenomenology of grounding, because it connects what an object is in the mind (a pattern-constellation) to how it is described (a linguistic pattern-constellation).

x^o and x^c can be seen as labels for $\{X\}$ and $\langle x \rangle$ respectively, and we must not forget that in physical and not functional terms, x^o and x^c are the same token x.

When we use a word in abstract conversation (x^c), we have the palpable sense it is “about” something in the world. That is because our cognitive system operates $\mathcal{R}(\{X\}, \langle x \rangle)$ — co-activating the experiential attractor $\{X\}$ and the linguistic pattern $\langle x \rangle$. The word’s apparent aboutness is the experiential shadow of an architectural coupling, what we call “grounding.”

8.4. What Grounding Looks Like

The formal characterization of R may suggest a static relation: {X} on one side, ⟨x⟩ on the other, R holding them together. But grounding is not a bridge built once at acquisition. It is a living coupling, continuously reshaped by encounter, because {X} is not a snapshot stored at the moment of learning. It is the current state of an attractor basin that has been sculpted by every encounter the organism has had with X, weighted by recency, affect, and salience.

Consider a speaker who has read extensively about black swans — Taleb’s metaphor, the falsification of “all swans are white,” the epistemology of improbable events. Her ⟨black swan⟩ is rich: inferentially connected, culturally loaded, deployable across registers. She can use “black swan” competently in any conversation. But her {SWAN}, the experiential attractor basin formed through encounters with actual swans, may contain no black swans at all. {SWAN} is white, serene, gliding — shaped by the swans she has actually seen.

Then she sees two black swans on a lake. The encounter is striking — the unexpected dark plumage against water, the size, the strangeness of a familiar form in an unfamiliar color. {SWAN} is resculpted. The attractor basin now integrates this visual-affective configuration alongside the previous one. Every subsequent activation of “black swan” co-activates it. The word has not changed. ⟨black swan⟩ has not changed. What has changed is {SWAN} — the experiential pole — and with it the character of R({SWAN}, ⟨swan⟩). The word now *means differently* for this speaker, not because its linguistic profile shifted but because what it is tethered to has been reshaped by the world; her *basic grounding* changed.

This is what distinguishes {X} from a representation. A representation has its content fixed at encoding, it is fundamentally a *template*. It can be copied, stored, and transmitted without alteration. {X} has no fixed content. It is the current state of a learning process that began with the organism's first encounters and continues for as long as the organism encounters. It is not *about* swans the way a mental file is about swans. It *is* the organism's current way-of-being-toward-swans — a dispositional landscape that the world continues to sculpt.

This is also why isomorphism theories of grounding capture something real but misidentify what they have captured. The geometric correspondence between ⟨x⟩-space relations and worldly relations is genuine — but it is inherited. It exists because the humans who produced the linguistic traces had {X}s that were shaped by encounter with the world. The isomorphism is a fossil record of past grounding: it preserves the *shape* of experiential structure without containing any. An LLM navigates this fossil record with statistical precision. It cannot add new encounters to it, because it has no {X}s to reshape. Its ⟨black swan⟩ is as rich as any speaker's. But nothing happens to it at the lake.

8.5. The Five Types of Grounding Revisited

The meta-language makes visible what the five types of grounding in (Mollo and Millière, 2025) actually are: different configurations of $\mathcal{R}$ with different argument types.

- **Referential grounding (classical)**: $\mathcal{R}(object, \langle x \rangle)$ — word hooks directly onto a worldly entity. This is what the standard debate seeks but cannot reach, because "object" in the formula is itself a linguistic description.
- **Sensorimotor grounding**: $\mathcal{R}(\{X\}, \langle x \rangle)$ — word coordinates with internal experiential pattern. This is closest to the PRU account, but the standard formulation treats $\{X\}$ as an "internal representation," collapsing it back into the linguistic register.
- **Relational grounding**: $\mathcal{R}(\langle x \rangle, \langle y \rangle)$ — word coordinates with word. The carousel in its purest form: both arguments are linguistic.
- **Communicative grounding**: Presupposes $\mathcal{R}$ is already in place and describes its social conditions.
- **Epistemic grounding**: Presupposes $\mathcal{R}$ is already in place and describes its evidential conditions.

The first four fail not because they are the wrong *kind* of relation, but because they have the wrong *arguments*. Relational grounding stays inside $\langle x \rangle$-space. Classical referential grounding tries to reach the bare object, which language cannot do. Communicative and epistemic grounding presuppose the coupling they need to explain. Sensorimotor grounding was closest, but needed the right characterization of what the experiential pole is — not a representation, but a constitutively subjective attractor basin formed through encounter.

8.6. Reference and Predication as $\mathcal{R}$-Configurations

Because $\mathcal{R}$ is a single operator, there is no need for a separate predication operator. What we call "predication" is $\mathcal{R}$ applied with different arguments. The configurations:

Reference: $\mathcal{R}(x^c, x^o)$. The operator takes the concept-word and the object-word for the *same* lexical base and produces the grounding of "x" — the tethering of the word's linguistic life to its experiential origin.

“What does ‘dog’ refer to?” is answered by the operation $\mathcal{R}(dog^c, dog^o)$.

Predication on an instance: $\mathcal{R}(p^c, x^o)$. The operator takes a property-concept and an object-word labeling an experienced instance. It produces a grounded predication: “This x here is p.” The property is applied to the instance through the same operation that produces reference — the only difference is that the two arguments have different lexical bases.

Kind-level predication: $\mathcal{R}(p^c, x^c)$. The operator takes a property-concept and a kind-concept. Both arguments are in the $\exists$-realm. It produces: “Xs are p.” This operates purely within linguistic space, though both p^c and x^c may carry residual grounding in their respective $\{P\}$ and $\{X\}$.

Baptism: $\mathcal{R}(name^c, h^o)$. The operator takes a name-concept and an object-word pointing to an experienced instance. It produces the initial grounding of a proper name: “This shall be called ‘Aristotle.’ ” $\mathcal{R}(Aristotle^c, h^o)$ is the first event that begins building $\mathcal{R}(\{ARISTOTLE\}, \langle aristotle\rangle)$.

8.7. Recursive Application

$\mathcal{R}$ supports recursive application in two directions:

Left-nesting (toward presence): $\mathcal{R}(\mathcal{R}(x^c, x^o), x^o)$. Each application reinforces x^o as an instantiator. The iterations contract toward $\{X\}$ — pulling the word back into the E-realm, toward immediate experiential presence.

Right-nesting (toward abstraction): $\mathcal{R}(x^c, \mathcal{R}(x^c, x^o))$. The concept-word keeps re-entering without reaching a final grounding in a particular instance. Semantic expansion without closure, remaining in the $\exists$-realm.

Language mediates between an E **pole** (x^o, sensory immediacy, $\{X\}$ activation, grounded experience) and an $\exists$ **pole** (x^c, conceptual elaboration, LPC operations, abstract thought). Each application of $\mathcal{R}$ can move toward either pole.

8.8. Why the Problem Looked Like a Problem

Traditional philosophy of language treats the word, the concept, and the referent as three separate entities requiring connection. Then it invents mechanisms — causal chains (Kripke), descriptions (Frege), mental representations (Fodor) — to bridge gaps between them.

PRU shows they are not three separate things. They are different aspects of the same learned cognitive structure, unified by a single operator:

- The word-token “x” as x^o: the experiential argument of $\mathcal{R}$.
- The word-token “x” as x^c: the linguistic argument of $\mathcal{R}$.
- The referent: what $\{X\}$ is an attractor *of* — the worldly encounters sedimented in the constellation.

The apparent mystery of reference is an artifact created by mistaking different arguments of the same operator for different entities, overlooking that experience and language co-form through learning, and seeking connections between things that were never separate.

The identity $\mathcal{R}(x^c, x^o) \equiv \mathcal{R}(\{X\}, \langle x\rangle)$ does not eliminate reference— it reframes it by showing why it appeared intractable under the standard framing.

8.9. Non-Commutativity and the Formal Nature of $\mathcal{R}$

That $\mathcal{R}$ is an operator and not a relation has formal consequences. An operator resolves a field of potentials into definite engagement — a word's potential for both x^o and x^c use resolves in context into a specific grounded configuration. This structure admits formalization as a projective operator in the quantum-mechanical sense, following the quantum cognition program that models cognitive states as vectors in Hilbert space and cognitive operations as projections (Pothos and Busemeyer, 2013), (Bruza and Busemeyer, 2012). Preliminary experimental work on non-commutative order effects in adjective-noun combinations supports this direction: predication behaves like measurement in that the order of operations matters — "the large red ball" and "the red large ball" yield measurably different cognitive states. $\mathcal{R}(large^c, \mathcal{R}(red^c, ball^o)) \neq \mathcal{R}(red^c, \mathcal{R}(large^c, ball^o))$.

The full development of this quantum-formal treatment is beyond the scope of the present paper and is explored separately in [Cojocariu, forthcoming].

9. From Generics to Predication

9.1. Extended Constellations

As children develop language, property-words enter constellations before separating into independent properties. When a child repeatedly hears "birds fly" while experiencing birds, fly^c integrates into the bird-constellation:

Stage 3.5 (extended constellation):

$$\{BIRD, bird^o, fly^c\}$$

The property-word becomes part of what birds are, not yet abstracted as independent. Formation:

$$(bird^o, fly^c) \to \{BIRD, bird^o, fly^c\} \to \{\{BIRD, bird^o\}, fly^c\}$$

This is how generics appear. They are more primitive than full predication.

9.2. Unified Pattern Formation

Notation: $\{\{X, x^o\}, p^c\}$

Reads as: "the unified object–property pattern."

Definition: A crystallized pattern that recognizes object and property as an inseparable gestalt. The double braces mark a pre-predicative unity: before analytical decomposition separates subject and predicate.

Why x^o, not x^c: Unified patterns are grounded in $E(\{X, x^o\})$ events where x^o labels the experienced instance. When infants say "bird" and point to one, they use $bird^o$. $bird^c$ appears at a later stage.

Examples:

- $\{\{BIRD, bird^o\}, fly^c\}$ — "birds fly" (kind-referring).

- $\{\{SNOW, snow^o\}, white^c\}$ — "snow is white" (mass noun, kind-referring).
- $\{\{TIGER, tiger^o\}, stripes^c\}$ — "tigers have stripes" (kind-referring).

Key distinction:

$$\{\{X, x^o\}, p^c\} \neq \mathcal{R}(p^c, x^c)$$

The first is unified pattern recognition — one event, at the PC level. The second is compositional predication: combining separate elements using the grounding operator.

9.3. Generic Sentences as Pattern Expression

Generic sentences express unified patterns.

"Birds fly":

$$\exists\,\{\{BIRD, bird^o\}, fly^c\}$$

This says: there is a crystallized bird–fly pattern built from many experiential episodes of the form $E(\{BIRD, bird^o, fly^c\})$.

It is not equivalent to "All birds fly" (universal quantification), nor to $\mathcal{R}(fly^c, bird^c)$ (a simple predicative proposition over two concepts).

This explains why generics appear before explicit quantifier predication developmentally (children say "birds fly" before they handle "all" and "some"), resist simple truth-conditional analysis, tolerate ex^ceptions (penguins do not falsify "birds fly"), and feel cognitively primitive — they express pattern recognition directly.

These features of generics are well documented in the empirical and philosophical literature. (Leslie, 2007) argues that generic generalizations are a cognitively default mode of predication, developmentally prior to quantified statements, resistant to truth-conditional reduction, and tolerant of ex^ceptions — properties that follow naturally if generics express unified patterns rather than quantificational structures. See also (Leslie, 2007) "Generics and the Structure of the Mind." *Philosophical Perspectives*, 21, 375–403.

9.4. Property Abstraction

Sometimes properties abstract across multiple constellations. A commonality is noticed across unified patterns:

- $\{\{BIRD, bird^o\}, fly^c\}$
- $\{\{PLANE, plane^o\}, fly^c\}$
- $\{\{BEE, bee^o\}, fly^c\}$

This abstraction enables a more general concept of flying, but it does not replace the unified patterns. Both structures coexist in mature cognition.

Critical asymmetry: Properties have no x^o mode in the sense of labeling experienceable objects. There is no object “fly” to encounter through $E(\{FLY\})$ events; there is always something else that flies. Properties are always abstractions across patterns, never grounded in the same way objects are. As Frege observed, properties are unsaturated. In PRU terms, they act as unsaturated $\mathcal{R}$s: $\mathcal{R}(\dots, x^c)$.

9.5. The Genesis of Predication

Predication emerges from decomposing unified patterns.

Developmental sequence:

1. $(bird^o, fly^c)$ — heard about an instance: “Bird flies.”
2. $\{BIRD, bird^o, fly^c\}$ — repeated instances build an extended constellation.
3. $\{\{BIRD, bird^o\}, fly^c\}$ — unified pattern crystallizes (generic “birds fly”).
4. $\langle fly\rangle, fly^c$ — property abstracts (sometimes).
5. $\mathcal{R}(fly^c, bird^o)$ — simple (natural) predication emerges (“This bird flies”).
6. $\mathcal{R}(fly^c, bird^c)$ — quantified predication emerges (“All birds fly”).

Children master step 3 (generics) and 5 (natural predication) before step 6 (quantified predication). Predication is the sophisticated achievement, not the foundation. Quantified predication does not enable pattern expression. It emerges from pattern recognition through a form of induction over experiential episodes.

10. Complete System Summary

10.1. Core Ontology

Experiential Realm (E):

- $E(\{A\})$: pre-linguistic pattern-recognition.
- $E(\{A, a^o\})$: experience with integrated label; a^o labels an instance.
- Sensory–motor–affective grounding.
- attractor states (of the kind modeled in Hopfield networks).

Linguistic Realm ($\exists$):

- $\exists(\{A, a^c\})$: conceptual mediation for A using a^c.
- $\exists(\langle a\rangle)$: activation of LPCs.
- Abstract, can operate without current E.

10.2. Developmental Hierarchy

1. $\{A\}$ — primitive PC, internally grounded (input-response unified).
2. $\{\{X\}, \{Y\}\}$ — composite PCs capturing relational structure.
3. Basic language — sounds inside PCs, triggering pattern completion in others.

4. $\{A, a^o\}$ — human label integration (continuity with basic language).
5. $\{A, a^o, a^c\}$ — concept-word emerges; word operates outside $\{A\}$.
6. $\langle a \rangle$ — LPC formed from patterns of a^c-use.
7. $\mathcal{R}(\{A\}, \langle a \rangle)$ — grounding operator tethers linguistic to experiential.

10.3. Pattern Hierarchy

Level	Element	Realm	Ground
Primitive PC	$\{A\}$	E	Internal (input-response unified)
Composite PC	$\{\{X\}, \{Y\}\}$	E	Situation-level integration
Labeled PC	$\{A, a^o\}$	E	Instance + word-sound
Full PC	$\{A, a^o, a^c\}$	E and $\exists$	Experience + concept
Unified pattern	$\{\{X, x^o\}, p^c\}$	E	Generic (pre-predicative)
LPC	$\langle a \rangle$	$\exists$	Community usage

10.4. The Grounding Operator $\mathcal{R}$

Level	Form	What it produces
Word-level (surface)	$\mathcal{R}(x^c, x^o)$	Grounded reference
Architectural (deep)	$\mathcal{R}(\{X\}, \langle x \rangle)$	Experiential-linguistic grounding

Fundamental identity: $\mathcal{R}(x^c, x^o) \equiv \mathcal{R}(\{X\}, \langle x \rangle)$

$\mathcal{R}$ appears where language creates separation. In basic minds, there is no $\mathcal{R}$ because there is no separation to bridge — grounding is internal to PCs.

10.5. Constitutive Asymmetries

	$\{X\}$	$\langle x \rangle$
Formation	Individual encounter	Collective usage
Access	First-person only	Publicly articulable
Transmission	Cannot be transmitted via text	Can be transmitted via text
Variation	Differs between individuals	Approximately shared
Character	Constitutively subjective	Constitutively collective

10.6. Notation Summary

Symbol	Meaning
$\{X\}$	Experiential Pattern-Constellation (attractor basin)
$\{\{X\}, \{Y\}\}$	Composite PC (situation-level integration)

Symbol	Meaning
$\langle x \rangle$	Linguistic Pattern-Constellation (pattern of word use)
x^o	Object-word mode (labels experienced instances)
x^c	Concept-word mode (operates in linguistic space)
$E(\{X\})$	Experiential activation of $\{X\}$
$E(\{X\}, x^o)$	Experiential activation with linguistic label
$\exists(\langle x \rangle)$	Linguistic mediation through concept-mode
$\mathcal{R}(x^c, x^o)$	Grounding operator, surface level
$\mathcal{R}(\{X\}, \langle x \rangle)$	Grounding operator, architectural level
$\{\{X, x^o\}, p^c\}$	Unified object–property pattern (generic)
$\exists\,\{\{X, x^o\}, p^c\}$	Generic sentence expressing unified pattern

11. The LLM Diagnostic

The meta-language developed in this paper makes precise what natural language habitually collapses: the difference between navigating $\langle x \rangle$ and being grounded in $\{X\}$. If the foregoing architecture is correct, then the meta-language yields the following diagnostic for current AI systems.

An LLM trained on a large corpus acquires $\langle x \rangle$ at high fidelity. Its internal geometry reflects the structure of human linguistic practice — the distributional patterns, the inferential connections, the discourse norms. In PRU terms, an LLM is a system with rich $\exists(\langle x \rangle)$ but no $E(\{X\})$. It navigates the linguistic realm comprehensively but has no experiential attractor basins formed through embodied encounter.

The grounding operator $\mathcal{R}(\{X\}, \langle x \rangle)$ cannot operate in such a system — not because of a missing component that could be added, but because $\mathcal{R}$ requires a structural coupling between experiential and linguistic capacities that co-developed through learning. In the LLM, there was never a stage at which word-sounds were *inside* experiential constellations — never a basic-language phase from which concept-words could emerge. The word "dog" in an LLM was never a motor-sensory pattern inside $\{DOG\}$ that then acquired an autonomous life in $<dog>$. It started and remained in $\langle\rangle$-space. $\mathcal{R}$ cannot operate because one of its arguments was never formed and the developmental history that enables the operation never occurred.

This returns us to the carousel diagnosed in §1. The standard grounding debate asks whether LLMs have grounding — but asks the question in language, using concepts about grounding, and evaluates the answer using linguistic criteria. LLMs, which are $\langle x \rangle$-type systems par ex^cellence, naturally satisfy $\langle x \rangle$-type conditions. The evaluation succeeds because evaluator and evaluated share the same register. The experiential pole — $\{X\}$ — is not detected as missing because the framework has no way to detect it.

The PRU notation does not solve this problem by magic. It does not deliver $\{X\}$ to the reader any more than natural language can. What it does is maintain *representational hygiene*: it prevents the specific conflation — of the collective linguistic pole with the full structure of grounding — from occurring unnoticed. With $\{X\}$ and $\langle x \rangle$ held apart by notation, the question of whether a system is grounded becomes: can $\mathcal{R}(\{X\}, \langle x \rangle)$ operate? And that question cannot be answered by inspecting $\langle x \rangle$ alone, no matter how rich or well-structured it is.

The difficulty is not merely theoretical. There is a systematic reason why the absence of {X} in LLMs is hard to detect, and it has to do with who is doing the detecting. The communities that build and evaluate AI systems — computer scientists, computational linguists, machine learning engineers, philosophers of mind — are overwhelmingly knowledge workers whose professional lives are spent navigating $\langle x \rangle$-space. Reading, writing, inferring, arguing, citing: these are operations in the linguistic realm. For such evaluators, $\langle x \rangle$-navigation *is* what competence feels like from the inside. The experience of understanding a topic, for a text-worker, largely *is* the experience of being able to navigate its linguistic pattern-constellation fluently.

(Bourdieu et al., 2000) identified this structural bias as the *scholastic fallacy*: the tendency of intellectuals to project the conditions of their own contemplative practice onto the nature of cognition itself, mistaking the logic of texts for the logic of things. The fallacy is especially acute in AI evaluation. When a researcher finds that an LLM navigates $\langle x \rangle$-space with fluency — answering questions, drawing inferences, producing well-formed arguments — the temptation to conclude that it *understands* is nearly irresistible, because fluent $\langle x \rangle$-navigation is what understanding looks like from inside a life conducted predominantly in language.

The meta-language makes this conflation visible. A system with $\exists(\langle x \rangle)$ but no $E(\{X\})$ is indistinguishable from a system where $R(\{X\}, \langle x \rangle)$ operates — *if the evaluation is conducted entirely within $\langle x \rangle$-space*. Every linguistic test, every benchmark, every conversational probe examines the system's $\langle x \rangle$-navigation. The experiential pole is not detected as missing because the evaluation has no instruments that reach it. This is not a contingent limitation awaiting a better benchmark. It is a structural consequence of evaluating grounding from within the register that cannot access it — the carousel, now operating not in the theory but in the methodology of AI research itself.

The full development of this diagnostic and its application to specific arguments in the grounding debate is given in [Cojocariu, forthcoming].

References

[1] Harnad S. The Symbol Grounding Problem 1990.

[2] Hodges W. Tarski's Truth Definitions. In: Zalta EN, Nodelman U, editors. The Stanford Encyclopedia of Philosophy. Winter 2022, Metaphysics Research Lab, Stanford University; 2022.

[3] Mollo DC, Millière R. The Vector Grounding Problem 2025. https://doi.org/10.48550/arXiv.2304.01481.

[4] Cole D. The Chinese Room Argument. In: Zalta EN, Nodelman U, editors. The Stanford Encyclopedia of Philosophy. Spring 2026, Metaphysics Research Lab, Stanford University; 2026.

[5] Varela FJ, Thompson E, Rosch E. The embodied mind: Cognitive science and human experience. revised edition. Cambridge, Massachusetts ; London England: MIT Press; 2016.

[6] Gibson JJ. The Ecological Approach to Visual Perception 2015.

[7] Hutto DD, Myin E. Radicalizing enactivism: Basic minds without content. Cambridge (Mass.): MIT press; 2013.

[8] Barsalou LW. Perceptual symbol systems. Behavioral and Brain Sciences 1999;22:577–660. https://doi.org/10.1017/S0140525X99002149.

[9] Kripke SA. Naming and necessity. Rev. and enl. ed., paperback ed., Repr. Oxford: Blackwell; 1990.

[10] Frege G. Sense and Reference. THE PHILOSOPHICAL REVIEW 1892.

[11] Field H. Tarski's Theory of Truth. The Journal of Philosophy 1972;69:347. https://doi.org/10.2307/2024879.

[12] Hopfield JJ. Neural networks and physical systems with emergent collective computational abilities. Proceedings of the National Academy of Sciences 1982;79:2554–8. https://doi.org/10.1073/pnas.79.8.2554.

[13] Davis KL, Montag C. Selected Principles of Pankseppian Affective Neuroscience. Frontiers in Neuroscience 2019;12. https://doi.org/10.3389/fnins.2018.01025.

[14] Montag C, Davis KL. Affective Neuroscience Theory and Personality: An Update. Personality Neuroscience 2018;1:e12. https://doi.org/10.1017/pen.2018.10.

[15] Seyfarth RM, Cheney DL, Marler P. Monkey responses to three different alarm calls: Evidence of predator classification and semantic communication. Science 1980;210:801–3. https://doi.org/10.1126/science.7433999.

[16] Tomasello M. Constructing a Language: A Usage-Based Theory of Language Acquisition. Cambridge: Harvard University Press; 2009.

[17] Bloom P. Précis of *How Children Learn the Meanings of Words*. Behavioral and Brain Sciences 2001;24:1095–103. https://doi.org/10.1017/S0140525X01000139.

[18] Carnap R. The logical structure of the world. Berkeley: Univ. of California Pr; 1969.

[19] Pothos EM, Busemeyer JR. Can quantum probability provide a new direction for cognitive modeling? Behavioral and Brain Sciences 2013;36:255–74. https://doi.org/10.1017/S0140525X12001525.

[20] Bruza PD, Busemeyer JR. Quantum Models of Cognition and Decision 2012.

[21] Leslie S-J. GENERICS AND THE STRUCTURE OF THE MIND. Philosophical Perspectives 2007;21:375–403. https://doi.org/10.1111/j.1520-8583.2007.00138.x.

[22] Bourdieu P, Nice R, Bourdieu P. Pascalian meditations. Stanford, Calif: Stanford University Press; 2000.