

Representation: When One Thing Stands For Another

IDE-012 · Fundamental Ideas

Concept essays for The Universe and Its Observers

Author: Juan Segurolo

Version: v0.9 EN · 19 September 2026

Status: Publication version

Article form: Concept essay

Scale / context: Maps, measurements, symbols, models and neural codes

Boundary: The stand-in relation itself. How much information a signal carries belongs to IDE-002; whether a model is adequate for a given inference belongs to IDE-015

Related works: IDE-013, IDE-002, IDE-015

Epistemic profile: Established knowledge

Licence at publication: All Rights Reserved

Article Abstract

A representation is something used to stand for something else. Four things are needed before the word does any work: a vehicle, the thing doing the standing; a target, the thing stood for; a mapping between distinctions in one and distinctions in the other; and a task the mapping is supposed to serve.

Representing does not require resembling, which is why a photograph, a differential equation and a pattern of neural activity can all represent the same object.[1,4] This article argues that representations are selective by construction rather than by failure, that statistical structure is not semantic content, and that representational language earns its keep only when the mapping can be specified rather than assumed. Whether a model is adequate for a given inference is left to IDE-015.

In 60 Seconds

Nothing represents anything by itself. A curve on a page becomes a representation only when someone uses it for a purpose,[4] which is why this article treats representation as a four-place relation, adding the mapping, rather than as a property of the vehicle. Once that is accepted, several confusions dissolve. Resemblance drops out: a map's usefulness comes from a systematic correspondence between its distinctions and the terrain's, not from looking like the ground, and every projection distorts something in order to preserve something else.[6] Selectivity stops being a defect: a representation that preserved everything would be the thing itself, so what matters is whether what was dropped was irrelevant to the task. And quantity is not content: Shannon's theory measures how much uncertainty a message resolves and set semantic aspects aside as irrelevant to the engineering problem,[3] so a high bit count says nothing about what a signal is about. Finally, representation is a hypothesis about a system, not a free description of it: attributing one to a mechanism means committing to a mapping that could be wrong.

Why This Question Matters

Representation is where several fields meet and talk past each other. Cognitive science asks whether a brain state represents an edge; metrology asks what a reading stands for; cartography asks what a projection preserves; machine learning asks what a learned vector encodes. All four are asking about the same relation, and all four go wrong in the same two ways: assuming a mapping without stating it, and treating the loss of detail as an imperfection rather than as the point. Getting the relation explicit is what makes it possible to say when a representation is wrong rather than merely different.

1. Four Things, Not One

The minimal analysis has four parts. There is a vehicle: marks on paper, a voltage, a data structure, a firing pattern. There is a target: the terrain, the temperature, the object in the visual field. There is a mapping that connects distinctions available in the vehicle to distinctions in the target. And there is a task, because a mapping adequate for one purpose can be useless for another.[4]

Giere's formulation makes the user and the purpose explicit rather than treating representation as a two-place relation between a thing and what it depicts.[4] That matters because the two-place version invites the question of what makes a mark intrinsically about something, which has no good answer. The four-place version replaces it with questions that can be settled: which distinctions map to which, and whether that mapping supports the inference being drawn.

Resemblance is then visibly beside the point. A wiring diagram does not look like a circuit; a differential equation does not look like a falling body; the word cat does not look like a cat. What these have in common is not similarity[4] but a systematic correspondence adequate to the task.

Where That Leaves Us

Representation needs a vehicle, a target, a mapping and a task, and stating all four turns an unanswerable question about intrinsic aboutness into answerable ones about correspondence and adequacy. Resemblance is neither required nor sufficient. What follows immediately is that every representation leaves something out, and that this is a design decision rather than a shortcoming.

2. Selective by Construction

A representation that preserved every distinction of its target would be a duplicate of it and would be no more tractable than the original. Selectivity is therefore constitutive: something must be dropped for the representation to be useful at all, and the interesting question is never whether detail was lost but whether the lost detail mattered for the task.

Map projections make this concrete because distortion is unavoidable rather than merely practical. Flattening a three-dimensional world onto a flat map necessarily distorts it. The choice of projection is therefore a choice about what the map is for. Monmonier's point is that this makes every map an argument, and that the arguments are usually invisible to readers who take the projection as given.[6]

The same structure appears wherever a representation is chosen. A summary statistic drops the distribution; a schematic drops the geometry; a categorical label drops the gradient. Each of those is defensible for some inferences and misleading for others, and the way to keep track is to name the task before judging the representation.

Where That Leaves Us

Every representation omits, because one that omitted nothing would be a copy, so the question is whether the omission is harmless for the task at hand. Map projections show the trade-off directly: flattening a three-dimensional world onto a flat map necessarily distorts it. A related confusion concerns not what is dropped but what is kept, and specifically whether keeping structure amounts to keeping meaning.

3. Structure Is Not Content

Shannon opened his 1948 paper by setting semantic questions aside, on the grounds that the engineering problem of reproducing a message at another point is independent of whether the message means anything.[3] That was a deliberate restriction of scope and it succeeded spectacularly, but it means the resulting quantities do not measure content.

The consequences are easy to state and easy to forget. A random string can have maximal entropy and represent nothing. A short message can be highly meaningful to a receiver holding the right convention and meaningless to one who does not. Two files can have identical statistics and different subject matter. None of this is a limitation of Shannon's theory; it is the boundary the theory drew for itself.

For representation, the moral is that a statistical relationship between a vehicle and a target is not sufficient. On the functional interpretation used here, cognitive representations are parts of a system's functional or computational organization, not merely correlations assigned by an outside observer.[5] A representation can also be caused by something it does not represent, and can represent something that has not caused it.[7]

Where That Leaves Us

Shannon's measures were built to work without semantics, so statistical structure in a vehicle does not by itself make it about anything. That raises the question of at which level of description a representation is supposed to be identified.

4. Levels, and Why Implementation Does Not Fix Content

Marr argued that an information-processing system must be understood at three levels: what problem is being solved and why, what representation and algorithm are used, and how these are physically realized.[1] The levels are related but not reducible to one another, and a complete story at the bottom does not deliver the top.

That asymmetry has a direct consequence. Knowing every voltage in a circuit or every spike in a network tells you what the system does physically without telling you what it is computing, because the same physical process can implement different computations under different mappings. Content is fixed at the level where the mapping is specified, not at the level where the mechanism runs.

This is also why the same computation can run on radically different hardware, and why arguments about whether a system really represents something cannot be settled by looking more closely at the substrate. They are settled, if at all, by showing which mapping the system's own operation relies on.

Where That Leaves Us

Marr's levels separate the problem, the representation and algorithm, and the physical realization, and a complete account of the mechanism does not fix what is being computed. Content lives where the mapping is specified. That leaves a practical question: when is it worth positing a representation at all, rather than describing the mechanism directly?

5. When Representation Talk Earns Its Keep

Newell and Simon proposed that a physical symbol system, a system with structures that designate and processes that operate on them, has the necessary and sufficient means for general intelligent action, and offered this as an empirical hypothesis rather than a definition.[2] Framing it as empirical is the useful part: it makes representational vocabulary something that can be right or wrong about a given system.

Representational and computational descriptions gain empirical content when they explain cognitive organization beyond a merely physical-level description.[5]

Two failures follow from ignoring the test. Over-attribution reads representations into every correlation, which makes the notion vacuous. Under-attribution refuses them wherever a mechanism can be described, which loses the generalization that different mechanisms can implement the same mapping. Naming the vehicle, target, mapping and task avoids both, and it is the same discipline the article opened with.

Where That Leaves Us

Representational and computational descriptions gain empirical content when they explain cognitive organization beyond a merely physical-level description. Over- and under-attribution are both avoided by naming the four parts. IDE-013 takes up a target that representations organize constantly and rarely examine: time itself.

What We Know, and How

- This article analyses representation as a relation among a vehicle, a target, a mapping and a task; making the user and the purpose explicit is Giere's move, and it replaces unanswerable questions about intrinsic aboutness with answerable ones.[4]
- Selectivity is constitutive rather than a defect; map projections make the trade-off unavoidable, because flattening a three-dimensional world onto a flat map necessarily distorts it.[6]
- Shannon set the semantic aspects of messages aside as irrelevant to the engineering problem,[3] so statistical structure in a vehicle does not by itself establish what it is about.
- A complete account of physical implementation does not determine what is being computed, because the same process can realize different computations under different mappings.

What We Still Don't Know

- What criteria reliably distinguish a genuine internal representation from a correlation that the system does not itself use, in biological systems where the downstream use is hard to trace.
- How to characterize the content of learned distributed representations in artificial systems, where the mapping is not designed and may not decompose into interpretable parts.

- Whether representational vocabulary is indispensable for explaining any biological capacity, or whether it is always a convenient summary of a mechanism that could be described directly.

What Could Change Our Minds

- A demonstration that two physical systems with identical internal dynamics and identical inputs and outputs must be assigned the same computational content would collapse Marr's level distinction and make implementation sufficient to fix content.[1]
- An operational criterion that identified representational content from vehicle statistics alone, without reference to downstream use or task, and that survived cases where the same statistics serve different functions, would overturn the four-place analysis.[3,4]
- A well-characterized biological system in which behaviour depends on a stand-in for an absent object with no identifiable mapping between vehicle and target distinctions would show the analysis to be incomplete rather than merely hard to apply.[5]

If You Remember Only One Thing

A representation is a vehicle used, under a mapping, for a task; without naming those, calling something a representation describes nothing.

Related Articles

- IDE-013 — Time
- IDE-002 — Information
- IDE-015 — Model

References

[1] David Marr. Vision: A Computational Investigation into the Human Representation and Processing of Visual Information. W. H. Freeman (1982), chapter 1.

[2] Allen Newell; Herbert A. Simon. Computer Science as Empirical Inquiry: Symbols and Search. Communications of the ACM 19(3), 113-126 (1976). DOI 10.1145/360018.360022.

[3] Claude E. Shannon. A Mathematical Theory of Communication. Bell System Technical Journal 27, 379-423 and 623-656 (1948). DOI 10.1002/j.1538-7305.1948.tb01338.x; 10.1002/j.1538-7305.1948.tb00917.x.

[4] Ronald N. Giere. How Models Are Used to Represent Reality. Philosophy of Science 71(5), 742-752 (2004). DOI 10.1086/425063.

[5] Zenon W. Pylyshyn. Computation and Cognition: Toward a Foundation for Cognitive Science. MIT Press (1984).

[6] Mark Monmonier. How to Lie with Maps, 3rd ed. University of Chicago Press (2018).

[7] David Pitt. Mental Representation. The Stanford Encyclopedia of Philosophy, Spring 2020 edition, Edward N. Zalta (ed.), first published 30 March 2000, substantive revision 21 January 2020.

<https://plato.stanford.edu/entries/mental-representation/>