

**MIT Press • Open Encyclopedia of Cognitive Science**

# **Mental Representation**

**Manolo Martínez<sup>1</sup>**

<sup>1</sup>Universitat de Barcelona

**MIT Press**

**Published on:** Apr 28, 2025

**DOI:** <https://doi.org/10.21428/e2759450.b3e7ef9b>

**License:** [Creative Commons Attribution 4.0 International License \(CC-BY 4.0\)](https://creativecommons.org/licenses/by/4.0/)

Cognitive scientists often try to provide explanations of how a human, animal, or artificial cognitive system manages to perform a certain cognitive task. For this, they often find it useful to postulate the existence of entities, localized somewhere within that cognitive system, that stand for, or are about others, often localized outside of it. These are the entities that are called *mental representations*. Explanations of cognitive function try to work out how computations over these representations may result in the successful performance of a certain task. Mental representations are a central, if not uncontested, theoretical posit in cognitive science.

## History

More than two millennia ago, Aristotle already suggested that we think of the way that senses take in the external world as analogous to how “a piece of wax takes on the imprint of the signet ring without the iron or gold” ([Aristotle, 350 BCE/1987](#)). This is not so different from the kind of strategy that many contemporary cognitive scientists deploy when investigating how some task is performed: they try to identify some part of the cognitive system (Aristotle’s piece of wax) that somehow tracks the qualities of some external entity, relevant to the task (Aristotle’s signet), in the hope that they will then be in a position to explain how manipulation of the inner surrogate results in behavior relevant to the external entity. This surrogate is what is called a *mental representation* of the external entity.

Contemporary discussion of mental representations is often taken to start with Franz Brentano’s *Psychology from an Empirical Standpoint*. Brentano, an Austrian psychologist, famously claimed that the distinguishing feature of mental phenomena, what makes them “mental,” is that they involve “direction toward an object” ([Brentano, 1874/1973](#), p. 68). Suppose that a cab is approaching you and that you see it. Seeing it allows you to make your behavior relevant to the cab: you will get out of its way, perhaps, or hail it if it you need a cab. How is it possible that a cab all the way out there makes it possible for you, in here, to form a good (adaptive, etc.) behavioral plan? One popular answer is that you are in a mental state that has the approaching cab as its content: you have a *mental representation* of the approaching cab. But this does not completely dispel the mystery: the cab is out there in the world, and your mental state is somewhere within the confines of your body, after all. More needs to be said about how the latter manages to be about the former. Much of the recent history of the debate on mental representations can be seen as trying to cash out Aristotle’s signet ring metaphor: what exactly is transmitted from the object of representation to the representation, and how. In the last 40 years, a prominent idea has been that representations must be understood in terms of their biological function. This approach is known as *teleosemantics* (from the Greek *télos*, meaning “end” or “goal”). Although this idea had been anticipated by a few theorists ([Stampe, 1977](#) among them), the first full formulation of teleosemantics is Ruth Millikan’s immensely influential *Language, Thought, and Other Biological Categories* ([Millikan, 1984](#)).

A second, partially overlapping set of discussions that have a bearing on mental representations has to do with devising systems for the more or less mechanical generation of thoughts, starting from a basic vocabulary and a

set of transformation rules. Such interest can be found as early as Ramon Llull in the 13th century. Later, Descartes and Leibniz, among many others, made important contributions to this program. For the contemporary study of mental representations, Frege was influential with his *Begriffsschrift*, a proposal for a formal calculus of thought processes ([Frege, 1879/1970](#), p. 14).

Starting in the 20th century, these ideas have been put to various uses. Most important is perhaps the foundational cognitive scientific project: explaining cognitive function as, at least partly, the processing of and computation over mental representations ([Newell & Simon, 1976](#); [Marr, 1982/2010](#)). Most recently, philosophers and neuroscientists have become interested in bridging the gap between *mental* and *neural* representation: that is to say, between the means, broadly conceived, by which minds take in, think about, and act back on the world; and the way in which neural activity comes to carry information about the world and transforms it in ways that are increasingly relevant to the production of behavior ([Piccinini, 2020](#)).

## Core concepts

### *Representational content*

Representations are parts of a cognitive system that stand for or are about other entities, often (but not always) outside the cognitive system. That which a representation represents is typically called its *content*. Having content is often taken to be the defining property of representations. Content is generally used as an umbrella term for any semantic (meaning-related) properties a mental state may have (cf. [Cummins, 1989](#), p. 12)

Many accounts of mental content focus on *propositional* content: content that corresponds to (putative) states of affairs or events. For example, a propositional content might be that it is raining here now, that a cab is approaching, that the Earth is flat, that space is flat, and, in general, that *p*. Concepts (a paradigmatic example of representation), on the other hand, have *subpropositional* contents: the concept dog has dogs (or the property of being a dog) as content; the concept green, greenness (or the property of being green); the (singular) concept bad bunny has the Puerto Rican singer Bad Bunny as content. In general, concepts have as content entities that can in some sense be part of a proposition.

What concepts are is a matter of much debate ([Machery, 2009](#)) [see [Concepts](#)]. Still, many theorists believe that a sufficient condition for a certain individual to possess a concept is the *generality constraint* ([Evans, 1982](#), p. 104): if a system is able to represent (*conceptual*) contents of the form *a* is *F* and *b* is *G*, then it must also be able to represent *a* as *G* or *b* as *F* (cf. [Toribio, 2007](#), p. 446). Examples of *nonconceptual* content often include perceptual experience (which appears to be impossible to capture with discursive thought, just like an image cannot be expressed in a thousand words) or the contents of subpersonal representations (e.g., those in the early visual cortex).

In contemporary accounts of representation, it is common to distinguish between two kinds of content ([Millikan, 1984](#)). First, there is *indicative* content, which is in the business of informing its recipient. Indicative

(or descriptive) content is by far the most commonly discussed and studied, but most theorists also recognize an *imperative* kind of content that is in the business of getting its recipient to see to it that its content happens. Representations with both indicative and imperative contents are often called *pushmi-pullyu representations* ([Millikan, 1995](#)).

## Vehicles

The characterization of representation given here mentions “parts of a cognitive system.” This characterization is purposely vague. Entities as different as events (e.g., neural population activity), objects (e.g., single neurons), mechanisms (e.g., synchronized neural activity), and other perhaps more exotic entities (e.g., the external world, as a result of cognitive offloading or some similar process; [Beer & Williams, 2015](#)) have been postulated as bearers of representational content.

Sometimes one needs to talk about these entities without taking a stand on which contents they carry or even if they carry content at all. In such circumstances one calls them *vehicles*: (potential) carriers of representational content. Vehicles are usually individuated by how they affect cognitive processing downstream ([Drayson, 2018](#)).

## Format

The question about which kinds of computations are made possible by which kinds of vehicles has often been formulated as being about the *format* of representations. For example, a historically important controversy concerned whether the representations underlying mental imagery have symbolic structure, similar to sentences in a language (sometimes called *linguistic* or *propositional*; [Pylyshyn, 2020](#)), or rather are image-like (*iconic* or *imagistic*; [Kosslyn, 1996](#)). Although this particular imagery debate has become less central, the question of what representational formats are, and how they affect cognitive processing, has gained prominence in recent years.

One traditional approach to format relies on analogies with public vehicles of representation: iconic formats being thought of as picture-like and propositional formats as language-like. On the other hand, some widely recognized formats are less amenable to this kind of analogy, for example, connectionist architectures or conceptual spaces ([Gärdenfors, 2000](#)). While particular representational formats are often appealed to in cognitive science ([Larkin & Simon, 1987](#)) there are not many theoretical elucidations of format in general (but see [Vorms, 2011](#); [Coelho Mollo & Vernazzani, 2024](#)).

Many important distinctions about representations, such as analog versus digital or discrete versus continuous, are best conceived as distinctions about representational format ([Maley, 2011](#)).

## Structure

Many of the representations discussed in cognitive science and neuroscience (e.g., hippocampal cognitive maps) are entities whose structure mirrors that of a target phenomenon ([Poldrack, 2020](#)). These so-called structural representations have become an important topic of discussion ([Gładziejewski & Miłkowski, 2017](#); [Artiga, 2023](#)).

## Basic constraints on theories of representation

These are various widely agreed constraints that theories of representations should respect:

- Representation and represented roles are asymmetric. That is to say, if A represents B, it does not follow that B represents A.
- Theories of representation need to make room for *inapt* representational contents. The most commonly discussed form of inaptness is misrepresentation: many representations purport to present some entity or state of affairs faithfully, yet fail to do so. Scrub jays misled about the location of food caches or humans believing that a certain white wall is actually red provide examples of misrepresentation. Some representations are not in the business of informing but of *moving* its receiver to do something. These imperative representations can also be inapt, not by misrepresenting anything but by being commands that should not be obeyed or advice that should not be heeded.
- It should be possible for representations to have nonexistent entities as part of their content. Hallucination (as opposed to illusion) is a prominent example; beliefs involving, for example, tree spirits, are (in all likelihood) another.
- It should also be possible for representations to have *distal* (as opposed to proximal) entities as part of their content. For example, many neuroscientists would like to say that at least some neural activity in brain area V4 represents the color of (extramental) stimuli. A theory of representation that attributes to activity in V4 contents involving some concrete activity in, say, V1, as opposed to contents involving the yellowness of some cab, would not meet this distal contents constraint (see [Schulte, 2018](#)).

## Questions, controversies, and new developments

### *The representational status question and the content question*

Current debate on mental representations is often organized in terms of two interrelated but different questions. Although some theorists believe that both need to be tackled in tandem ([Shea, 2018](#)), many see them as relatively independent.

First, the *representational status* question: What makes some entities (and not others) representational vehicles at all? There are entities, processes, and properties (such as neural activity in V1, hippocampal maps, concepts, or corollary discharges) that many cognitive scientists are willing to claim are representations and very many

that they do not, such as, say, the state of individual ion channels or the color of neurons. We would like to have a principled way to tell them apart.

Second, the *content* question: Given that a certain entity is a representation, what should one take its content to be? Many researchers believe that place cell activity (at least sometimes) represents places and not, say, days of the week. We aim at devising principled ways of attributing representations with contents.

## ***Liberality***

There is also broad agreement that answers to the representational status question should not be unduly liberal, or unduly restrictive, in the kinds of entities they grant representational status to. Many theorists agree about the extreme, paradigmatic cases (e.g., river pebbles are not representations, and some of the neural activity in the ventral visual pathway probably is), but there is disagreement regarding the large gray area between those extremes. For example, cognitive scientists disagree about whether alarm calls in vervet monkeys are representational, what can be said about quorum sensing in bacteria, and about the status of neural activity in early visual areas. According to *liberal representationalism*, many of these controversial cases are representations ([Artiga, 2016](#), p. 409). Other theorists hold a more restrictive view, for instance, claiming that representational contents must be shown to be explanatorily useful before they are accepted in accounts of cognitive function ([Rescorla, 2013](#), p. 93).

## ***Theories of representation***

For a first look into the kinds of problems that crop up when trying to come up with answers to the representational status and content questions, consider this toy causation-based response to the content question (which no one has ever endorsed): a representation R has the content that C if and only if C causes R.

Take a certain candidate representation, a photograph. Consider what the particular distribution of colors across its surface represent. It is at least initially plausible that they represent whatever visible features of whatever object caused that particular distribution of colors. This intuition is vindicated by the above toy causal account. On the other hand, according to this account, a representation has as content whatever caused it to exist, and this means that it cannot be false; if a normally illuminated red wall and a white wall illuminated with a cleverly concealed red light bulb both cause a certain representation R, then R means something like “this wall is either red or white and illuminated with a red light bulb.” This is sometimes called the *problem of error* ([Ryder, 2019](#)) and sometimes the *disjunction problem* ([Fodor, 1990](#)) because of the disjunctive “A or B or ...” that these content attributions take.

This account is too simplistic, but it captures a central intuition of more sophisticated accounts: one needs some way to cash out the idea that representations and their objects are somehow connected. Consider again a cab approaching and the subsequent, perceptually based judgment that a cab is approaching. In which way are cab out there and judgment in here linked together? The notion that representations need to be caused by its object

is one way, but there are other ways to ensure this kind of linkage. One important lesson of the above toy causal account is that linkage does not ensure aptness; causal relations, informational relations, and other plausible candidates for linkage are not discriminating. Some extra ingredient is needed to filter out good from bad content candidates.

One of the first attempts at doing this in the contemporary (naturalistic) study of mental representation is in terms of *fidelity conditions* ([Stampe, 1977](#)). For instance, consider the number of rings in a tree stump that can be used as a representation of the number of years that the tree lived. Here, representation (number of rings) and content (number of years) are adequately connected only if the climatic conditions that prevailed during the growth of the tree are normal with no severe climatic stress, but if the tree had to pull through a few severe droughts and fires, then it might have fewer rings. Normal climatic conditions are the fidelity conditions for this representation.

One problem is that these fidelity conditions seem to be plucked out of thin air. One promise of causal accounts of representation is that one will be able to cash out what it is for something to represent something else in naturalistic terms. However, for example, it is unclear according to which fidelity conditions, decided by whom, activity in V4 represents the color of stimuli. If it is cognitive scientists that decide which fidelity conditions are relevant to any particular representation, in a sense, this is only a theory of what activity in V4 represents for cognitive scientists (cf. [Egan, 2025](#)), and this is not the kind of theory that one might wish for. What is needed is a theory of what it represents for the cognitive system that hosts such activity.

Much contemporary work on representation can be seen as attempting to provide naturalistically acceptable characterizations of the idea of fidelity conditions. Most of it takes place within the so-called *teleosemantic* program, by far the most popular, most actively developed research program in this area. Teleosemantics can be summarized as the view that what fixes the fidelity conditions of a representation is the *biological function* of the representational vehicle or mechanism. Hearts pump blood and make thumping noises. Both are useful effects (the latter, for example, for diagnostics), but the heart is for the former, not the latter. Hearts are supposed to pump blood, and when they do not, it is usually said that they malfunction. Many theorists, following [Millikan \(1984\)](#), call *proper functions*, or *teleofunctions*, those effects of designed or biological devices that underlie malfunction claims.

Consider the tree example. There is something awkward about thinking of tree rings as representations. Mere onlookers can rely on them to learn about the tree's age, but nothing in the tree itself reads those rings. They represent nothing to anything in the tree itself. It has been argued that representations are necessarily for something or someone; representations mediate between a *producer* that brings it into existence and a *consumer* that puts it to use (see [Figure 1](#)).

The exact same motif (upstream goes on, a producer, a representation, a consumer, and downstream consequences) is also relied on in the signaling framework ([Skyrms, 2010](#)), in which producers and consumers

are called *senders* and *receivers*, and again in information theory ([Shannon & Weaver, 1949/1998](#)), in which they are often called, respectively, *transmitters* or *encoders* and *decoders* (cf. [Cao, 2012](#)).



**Figure 1**


Schematic showing that representations are intermediate between producers and consumers of the representation.

The teleosemantic response to the representational status question is that a representation is an entity mediating between a producer and a consumer, in which both producer and consumer are mechanisms individuated by their proper function. In the mainstream *consumer* development of teleosemantics ([Millikan, 1984](#); [Papineau, 1987](#)), it is the consumer of the representation that fixes its content. Other prominent approaches aim instead at deriving content from causal informational ties between the representation and its producer. On one version of this idea, a representational system is one that has the function of indicating the states of some extramental phenomenon ([Dretske, 1988](#), p. 53). Here, the key concepts are *function* (which should be understood as meaning proper function) and *indication*, which is a probabilistic notion: A indicates B if the probability of B is lower than the probability of B conditional on A. A *representational system* is what teleosemanticists call a producer (sender, encoder, etc.); it has the function of producing a signal, or entering a state, that correlates with the indicated state of affairs (cf. [Neander, 2017](#), p. 151).

More recently, some theorists have offered what can be seen as more or less purely information theoretic accounts of representation. Recent work on so-called *signaling games* has aimed at substituting proper functions with game theoretic equilibria or dynamically stable configurations of sender and receiver populations in evolutionary game theoretic models ([Skyrms, 2010](#)).

The sender–receiver motif is one of those ideas that keeps being rediscovered. One of its early appearances, and perhaps the most influential of all, is in the point-to-point model of information transmission in [Shannon & Weaver \(1949/1998\)](#) (see [Figure 2](#)).



**Figure 2**

The Shannon and Weaver model for information transmission.

Some theorists aim at leveraging the full Shannon information transmission model in theories of representation ([Mann, 2023](#)).

### ***Are representations a useful theoretical posit?***

Mainstream cognitive science relies on representations and computations over them, as central theoretical posits in explanations of cognitive function; however, an increasingly influential family of approaches (sometimes collectively called *4E cognitive science*) disputes this centrality. For example, *embodied cognitive science* is the idea that cognition happens both in the brain and the body: cognitive agents act in ways that radically diminish the need for internal information processing ([Chemero, 2011](#)). For example, the task of catching a ball that someone else has thrown might be described as working out (computing) where it will land, given an estimation of its initial position and acceleration, and moving there or, alternatively, might be solved by running in a way so as to make the ball appear to move with constant velocity ([Wilson & Golonka, 2013](#), p. 5). In the second explanation, computations and representations play no meaningful role. They are replaced by clever behavior (the running alongside the ball) that keeps the complexity of the task low all the way through completion. An increasingly important question in contemporary cognitive science is how pervasive this kind of nonrepresentational cognitive strategy is. It has been argued that cognitive science is moving in this nonrepresentationalist direction ([Ramsey, 2007](#)).

### **Broader connections**

Mental representations are involved in much, perhaps most, research in cognitive science, and many of the central theoretical posits in those research efforts just are representations [see [Face Perception](#), [Theory of Mind](#), and [Concepts](#)]. The main source of evidence in favor of or against theories of representation comes from the way representations are relied on in cognitive science research.

It is also very common for theories in cognitive science, which do not necessarily take a stand on the representational status of their theoretical posits, to be interrogated as regards that status; a prominent recent example is work on the Free Energy Principle, with theorists arguing whether the principle should be read representationally (e.g., [Figdor, 2020](#)) or nonrepresentationally (e.g., [Bruineberg et al., 2018](#)) [see [Free Energy Principle](#)].

While most philosophical research on mental representations is qualitative, and makes little use of formal tools, there are theoretical projects that do; game theoretic work on signaling has established a fruitful dialogue with theories of mental representation [see [Signaling](#)]. Bayesian models of cognition are also explicitly representational [see [Bayesian Models of Cognition](#)].

There are various approaches in computational neuroscience to unveiling and understanding representation in the brain: encoding and decoding models, representational similarity analysis ([Kriegeskorte & Douglas, 2019](#)), or analyses based on deep learning ([Richards et al., 2019](#)). There has been philosophical work on these approaches, but more cross-talk is desirable ([Roskies, 2021](#); [Subotić, 2024](#)).

Finally, an important research program in the philosophical study of consciousness is *representationalism*: the idea that the phenomenal character of conscious states depends on, or perhaps is identical to, representational content ([Tye, 2000](#)). Glossing over the many different representationalist approaches to consciousness, this research program proposes that representational roles (and not, say, causal roles, synchronized activity, or firing patterns) provide the optimal “fineness of grain” on cognitive function for the purposes of understanding consciousness.

On the other hand, the main claim in the *phenomenal intentionality* research program is that the representational properties of mental states are grounded on their phenomenal profile: because they are conscious, they are representational ([Mendelovici, 2018](#)). Representationalism and phenomenal intentionality seem incompatible approaches to consciousness. Which one (if any) should carry the day or whether they could be, after all, reconciled are important open questions.

## Further reading

- Kriegeskorte, N., & Diedrichsen, J. (2019). Peeling the onion of brain representations. *Annual Review of Neuroscience*, 42, 407–432. <https://doi.org/10.1146/annurev-neuro-080317-061906>
- Millikan, R. G. (1984). *Language, thought and other biological categories*. MIT Press.
- Schulte, P. (2023). *Mental content*. Cambridge University Press.
- Shea, N. (2018). *Representation in cognitive science*. Oxford University Press.

## References

- Aristotle. (1987). *De anima (On the soul)* (H. Lawson-Tancred, Trans.). Penguin Classics. (Original work published 350 BCE.)

[↑](#)

- Artiga, M. (2016). Liberal representationalism: A deflationist defense. *Dialectica*, 70(3), 407–430. <https://doi.org/10.1111/1746-8361.12152>

[↑](#)

- Artiga, M. (2023) [Forthcoming]. Understanding structural representations. *The British Journal for the Philosophy of Science*. <https://doi.org/10.1086/728714>

↵

- Beer, R. D., & Williams, P. L. (2015). Information processing and dynamics in minimally cognitive agents. *Cognitive Science*, 39(1), 1–38. <https://doi.org/10.1111/cogs.12142>

↵

- Brentano, F. (1973). *Psychology from an empirical standpoint* (A. C. Rancurello & D. B. Terrell, Trans., L. L. McAlister, Ed. & Trans.). Routledge & Kegan Paul. (Original work published 1874)

↵

- Bruineberg, J., Kiverstein, J., & Rietveld, E. (2018). The anticipating brain is not a scientist: The free-energy principle from an ecological-enactive perspective. *Synthese*, 195, 2417–2444. <https://doi.org/10.1007/s11229-016-1239-1>

↵

- Cao, R. (2012). A teleosemantic approach to information in the brain. *Biology & Philosophy*, 27, 49–71. <https://doi.org/10.1007/s10539-011-9292-0>

↵

- Chemero, A. (2011). *Radical embodied cognitive science*. MIT Press.

↵

- Coelho Mollo, D., & Vernazzani, A. (2024). The formats of cognitive representation: A computational account. *Philosophy of Science*, 91(3), 682-701. <https://doi.org/10.1017/psa.2023.123>

↵

- Cummins, R. (1989). *Meaning and mental representation*. Bradford Books.

↵

- Drayson, Z. (2018). The realizers and vehicles of mental representation. *Studies in History and Philosophy of Science Part A*, 68, 80–87. <https://doi.org/10.1016/j.shpsa.2018.01.005>

↵

- Dretske, F. (1988). *Explaining behavior. Reasons in a world of causes*. MIT Press.

↵

- Egan, F. (2025). *Deflating mental representation*. The MIT Press.

↵

- Evans, G. (1982). *The varieties of reference* (J. McDowell, Ed.). Oxford University Press.

↵

- Figdor, C. (2020). Shannon + Friston = Content: Intentionality in predictive signaling systems. *Synthese*, 199, 2793–2816. <https://doi.org/10.1007/s11229-020-02912-9>

↵

- Fodor, J. A. (1990). *A theory of content and other essays*. MIT Press.

↵

- Frege, G. (1970). Begriffsschrift, a formula language, modeled upon that of arithmetic, for pure thought. In J. Van Heijenoort (Ed.), *Frege and Gödel* (pp. 1–82). Harvard University Press. (Original work published in 1879.)

↵

- Gärdenfors, P. (2000). *Conceptual spaces: The geometry of thought*. MIT Press.

↵

- Gładziejewski, P., & Miłkowski, M. (2017). Structural representations: Causally relevant and different from detectors. *Biology and Philosophy*, 32, 337–355. <https://doi.org/10.1007/s10539-017-9562-6>

↵

- Kosslyn, S. M. (1996). *Image and brain: The resolution of the imagery debate*. MIT Press.

↵

- Kriegeskorte, N., & Douglas, P. M. (2019). Interpreting encoding and decoding models. *Current Opinion in Neurobiology*, 55, 167–179. <https://doi.org/10.1016/j.conb.2019.04.002>

↵

- Larkin, J. H., & Simon, H. A. (1987). Why a diagram is (sometimes) worth ten thousand words. *Cognitive Science*, 11(1), 65–100. <https://doi.org/10.1111/j.1551-6708.1987.tb00863.x>

↵

- Machery, E. (2009). *Doing without concepts*. Oxford University Press.

↵

- Maley, C. J. (2011). Analog and digital, continuous and discrete. *Philosophical Studies*, 155, 117–131. <https://doi.org/10.1007/s11098-010-9562-8>

↵

-

[↑](#)

- Marr, D. (2010). *Vision: A computational investigation into the human representation and processing of visual information*. MIT Press. (Original work published in 1982.)

[↑](#)

- Mendelovici, A. A. (2018). *The phenomenal basis of intentionality*. Oxford University Press.

[↑](#)

- Millikan, R. G. (1984). *Language, thought, and other biological categories*. MIT Press.

[↑](#)

- Millikan, R. G. (1995). Pushmi-pullyu representations. *Philosophical Perspectives*, 9, 185–200.  
<https://doi.org/10.2307/2214217>

[↑](#)

- Neander, K. (2017). *A mark of the mental: In defense of informational teleosemantics*. MIT Press.

[↑](#)

- Newell, A., & Simon, H. A. (1976). Computer science as empirical inquiry: Symbols and search. *Communications of the ACM*, 19(3), 113–126. <https://doi.org/10.1145/360018.36002>

[↑](#)

- Papineau, D. (1987). *Reality and representation*. Basil Blackwell.

[↑](#)

- Piccinini, G. (2020). *Neurocognitive mechanisms: Explaining biological cognition*. Oxford University Press.

[↑](#)

- Poldrack, R. A. (2020). The physics of representation. *Synthese*, 199, 1307–1325.  
<https://doi.org/10.1007/s11229-020-02793-y>

[↑](#)

- Pylyshyn, Z. W. (2002). Mental imagery: In search of a theory. *Behavioral and Brain Sciences*, 25(2), 157–182. <https://doi.org/10.1017/s0140525x02000043>

[↑](#)

- Ramsey, W. M. (2007). *Representation reconsidered*. Cambridge University Press.

[↑](#)

- Rescorla, M. (2013). Millikan on honeybee navigation and communication. In D. Ryder, J. Kingsbury, & K. Williford (Eds.), *Millikan and her critics* (pp. 87–102). Wiley-Blackwell.

↑

- Richards, B. A., Lillicrap, T. P., Beaudoin, P., Bengio, W., Bogacz, R., Christensen, A., Clopath, C., Ponte Costa, R., de Berker, A., Ganguli, S., Gillon, C. J., Hafner, D., Kepecs, A., Kriegeskorte, N., Latham, P., Lindsay, G. W., Miller, K. D., Naud, R., Pack, C. C., ... Kording, K. P. (2019). A deep learning framework for neuroscience. *Nature Neuroscience*, 22, 1761–1770. <https://doi.org/10.1038/s41593-019-0520-2>

↑

- Roskies, A. L. (2021). Representational similarity analysis in neuroimaging: Proxy vehicles and provisional representations. *Synthese*, 199, 5917–5935. <https://doi.org/10.1007/s11229-021-03052-4>

↑

- Ryder, D. (2019). Problems of representation II: Naturalizing content. In S. Robbins, J. Symons, & P. Calvo (Eds.), *The Routledge companion to philosophy of psychology* (pp. 251-279). Routledge.

↑

- Schulte, P. (2018). Perceiving the world outside: How to solve the distality problem for informational teleosemantics. *The Philosophical Quarterly*, 68(271), 349–369. <https://doi.org/10.1093/pq/pqx052>

↑

- Shannon, C. E., & Weaver, W. (1998). *The mathematical theory of communication*. University of Illinois Press. ((Original work published 1949.)

↑

- Shea, N. (2018). *Representation in cognitive science*. Oxford University Press.

↑

- Skyrms, B. (2010). *Signals: Evolution, learning, and information*. Oxford University Press.

↑

- Stampe, D. W. (1977). Toward a causal theory of linguistic representation. *Midwest Studies in Philosophy*, 2(1), 42–63. <https://doi.org/10.1111/j.1475-4975.1977.tb00027.x>

↑

- Subotić, V. (2024). Exploring, expounding & ersatzing: A three-level account of deep learning models in cognitive neuroscience. *Synthese*, 203, 94. <https://doi.org/10.1007/s11229-024-04514-1>

↑

- Toribio, J. (2007). Nonconceptual content. *Philosophy Compass*, 2(3), 445–460. <https://doi.org/10.1111/j.1747-9991.2007.00075.x>

↑

- Tye, M. (2000). *Consciousness, color and content*. MIT Press.

[↵](#)

- Vorms, M. (2011). Formats of representation in scientific theorizing. In P. W. Humphreys & C. Imbert (Eds.), *Models, simulations, and representations* (pp. 250–273). Routledge.

[↵](#)

- Wilson, A. D. & Golonka, S. (2013). Embodied cognition is not what you think it is. *Frontiers in Psychology*, 4, 58. <https://doi.org/10.3389/fpsyg.2013.00058>

[↵](#)