

# Thinking of oneself as someone: the structure of minimal self-representation

Julian Hauser

## Abstract

When investigating the nature of self-representation, we might ask which types of property need to figure in their content. Here, authors have claimed that self-representations need to be about spatial, temporal, bodily, or mental properties. However, we can also ask a second question: *how* do we need to represent a property to *self*-represent it? It is this latter question that I address. I argue that a distinction between egocentric and non-egocentric forms of representation – known primarily from the literature on spatial cognition – also applies to representations of other kinds of property. I use examples drawn from animal cognition and developmental psychology to show how creatures non-egocentrically represent their temporal, bodily, and cognitive properties. These representations are, I submit, minimal self-representations: they involve representing one's properties so that an explicit differentiation is made between the system's and other objects' properties (or between the system's actual and merely possible properties), they are directly linked to behaviour and sensation, and they are immune to error through misidentification. The upshot is a view on which different creatures may self-represent (in this minimal sense) more or fewer kinds of property. More substantive forms of self-representation (for instance, as exemplified by neurotypical adult human beings) then require integrated minimal self-representations of the right kinds of property.

Representations may be about different kinds of property, and they may represent these properties in different ways. A red apple may be represented by an image of a red apple or the words 'red apple'. The matter is no different for cognitive (and mental) representations: I can token the phrase 'red apple' in inner speech or conjure up an image of a red apple. This paper concerns the special sorts of representation with which a system may represent its own properties. I argue that these representations are (minimal) *self*-representations when they represent properties in a specific way. Self-representation doesn't depend on *what* properties are represented but only on *how* they are represented.

A minimal self-representation is the simplest kind of representation that is a genuine self-representation. A genuine self-representation, a representation with *de se*

content, is a representation with which a system represents *itself qua itself*. To do so, it must fulfil two conditions. First, the representation must explicitly differentiate between properties exemplified by the system and those exemplified by the world. A creature that represents a *tree to be located to the left* doesn't self-represent since the fact that the object is to the left *of the creature itself* isn't made explicit. In other words, the representation must include a token that stands for the system and one that stands for the tree, rather than, say, a single token specifying the distance to a tree. In the latter case, the system may *use* the representational token as denoting a distance to itself, but that fact isn't explicitly represented. We say that such a representation merely *concerns* the self, whereas a self-representation is *about* the self (Perry, 1993). Explicitly differentiating between self and other doesn't ensure that the system knows what token refers to itself. An animal may look into a mirror, represent someone to have certain features, but fail to realise that it is that someone. This animal may represent itself but fail to do so *qua itself*.

Here, we arrive at the second condition: a self-representer must be disposed to update and use the represented information in a specific way that directly links it to sensation and behaviour (Evans, 1982). The creature must be disposed to use sensory information about the object in question to update its representation of itself. If someone were to draw a spot on the animal's forehead, it should represent itself as now having a spot on its forehead. And the animal must be disposed to use the represented information for behaviour in the relevant ways. For instance, the animal should infer that it needs to engage in certain motor actions to have an unstained forehead.

I argue that minimal self-representations are coordinated non-egocentric representations. We know non-egocentric (or allocentric) representations from the literature on spatial cognition, where they appear in the guise of cognitive maps (O'Keefe & Burgess, 1996; O'Keefe & Nadel, 1978; Rescorla, 2017). Like the maps on our phones, cognitive maps represent the locations of various objects relative to each other. A fruit bat, for instance, represents its nest and a fruit tree with two representational tokens (Tsoar et al., 2011). The relation between these tokens corresponds to the spatial relation between the two objects. Such representations contrast with egocentric representations such as those employed in path integration. The *Cataglyphis* desert ant, for instance, employs path integration to continuously update a representation of the *distance and bearing* to its nest (Collett, 2019; Gallistel, 1989; Reid et al., 2012; Srinivasan, 2015).

The issue of self-representation is linked to the striking distinctions between egocentric and non-egocentric spatial representations (see Grush, 2000). The token that the ant uses to represent the nest in the egocentric case specifies a relation between the nest and the ant's present position. However, the fact that this relation is represented does not figure in the representation's explicit content. Such a representation concerns the ant (its accuracy depends on the ant's location) but isn't about the ant (no representational token stands for the ant). The non-egocentric case is different:

a relation between representational tokens explicitly represents the spatial relation. With the help of a representational token that stands for the ant, such a representation explicitly differentiates between properties attributed to the system and those attributed to the rest of the world. This representational token is *about* the bat.

Another difference is striking, too: while the egocentric representation necessarily concerns the ant itself (always representing a relation *to the ant*), the representer may disappear in a non-egocentric representation. The bat may represent its cave to be such-and-such a distance away from a food source without locating itself on the cognitive map. Non-egocentric representations represent objects and their properties as in principle independent of the system's state. Hence, to use its non-egocentric representation for navigation, the bat must employ a specific representational token as specifying *its* location. When this happens, I call the representation *coordinated*. When the representation is coordinated, the system has a special sort of know-how regarding its non-egocentric representation. For instance, the bat knows how to calculate the motor commands required to reach its cave based on the represented relation between its location and the cave.

Spatial properties constitute merely one kind of property that organisms represent with coordinated non-egocentric representations. I discuss evidence to this effect for representations of temporal, bodily, and cognitive properties. For instance, the forward models postulated by certain computational approaches to motor control (Grush, 2004; Körding & Wolpert, 2006; Wolpert & Ghahramani, 2000) imply an explicit differentiation between various possible bodily states. This, I argue, means they are non-egocentric representations. To use such a model to predict one's future bodily state, the model needs to be supplied with information about the system's actual state, which makes it a coordinated representation.

Coordinated non-egocentric representations explicitly differentiate between the system's and other objects' properties or between the system's actual and merely possible properties. Coordination involves the system using a certain token in a special manner, namely as specifying its own state. Coordination entails, first, that relations from this token to other tokens imply sensorimotor information and that, hence, the system is disposed to use the represented information to guide behaviour. Second, coordination entails that sensorimotor information implies relations between this special token and other tokens so that the system is disposed to update the representation given relevant sensory input.

Since coordinated self-representations explicitly attribute properties to the self and link in the requisite way to behaviour and sensation, they are genuine (if minimal) self-representations. Whenever a system represents *any* property this way, it instantiates a minimal self-representation. As we will see in the examples I discuss, such self-representations are exemplified by infants and a range of non-human animals.

Moreover, the proposed account also explains why certain self-attributed properties are immune to errors through misidentification (Evans, 1982; Perry, 2010; Shoe-

maker, 1968; Wittgenstein, 2007). A self-representation that results from a system inferring its properties from sensorimotor information is immune in this sense as sensorimotor information is inherently self-concerned. Self-attributed properties that aren't immune to errors through misidentification are inferred by identifying oneself with some object – such my identification with the person I see reflected in the mirror in front of me.

I hope this paper proves interesting to researchers in the empirical sciences and those more theoretically inclined. Cognitive scientists may find in this paper a proposal of the cognitive processes involved in self-representation that can be operationalised by extending the well-known paradigms from the study of spatial cognition without thereby giving up philosophical rigour. The resulting studies should allow us to investigate if, when, and how human and non-human animals (and even artificial agents) exhibit self-representation. Self-representation – at least of the simpler varieties – will turn out to be more common than we might have thought. This, in turn, may prompt philosophers to broaden their investigations beyond the case of adult human beings that employ linguistic representations. Genuine (if minimal) self-representation doesn't require, as Musholt (2012) thinks, that we represent our mental states. Moreover, such representation isn't linked, in an essential way, to spatial cognition (Grush, 2000; Ismael, 2008) or the representation of our bodies (Hohwy & Michael, 2017; Metzinger, 2003), and it doesn't require the representation of temporal properties (Campbell, 1999; Peacocke, 2014).

I call accounts that require self-representations to carry specific kinds of content *substantive*. While my account suggests that substantive accounts of self-representation don't explain self-representation's deepest puzzles, they explain, for instance, the self-representations of neurotypical adult human beings. According to my account, properties must be represented non-egocentrically to figure in such a substantive self-representation. Moreover, the account suggests that many arguments in the philosophical literature may concern the kind of content needed to represent a self's essential properties. If this is right, then the disagreements between such accounts might concern the metaphysics of selfhood rather than questions of self-representation.

The paper is structured as follows. After this introduction follow four sections investigating how various creatures use coordinated non-egocentric representations. First, I discuss the representation of spatial properties and introduce the main concepts employed throughout the paper. This section, and the three sections that follow it, imply a link between coordinated non-egocentric representation and self-representation. In the fifth section, I argue that coordinated non-egocentric representations are genuine (if minimal) self-representations. Section 6, looks at how creatures infer their properties and when the resulting self-attributions are immune to certain errors through misidentification. Section seven links minimal self-representation to substantive self-representation. Section eight concludes.

## Spatial cognition and minimal self-representation

The distinction between egocentric and non-egocentric representation arose in the literature on spatial cognition, and it's here that some (Grush, 2000) have first suggested links between non-egocentric representations and self-representation. Additionally, path integration and cognitive maps provide a uniquely intuitive starting point, making spatial cognition the ideal first stop on the road to a general account of minimal self-representation. This section will distinguish between egocentric and non-egocentric representations in the spatial domain and give you a first taste of why coordinated non-egocentric representation is genuine self-representation.

The *Cataglyphis* desert ant has become a poster child for path integration (Collett, 2019; Gallistel, 1989; Reid et al., 2012; Srinivasan, 2015). These ants often leave their nests on lengthy foraging expeditions, and even though these trips can send them as far as 200 metres from home, the ants (almost) always manage to return. This achievement is particularly remarkable given that the ants' monotonous desert habitat is largely lacking in distinguishing features. Furthermore, if we capture a homeward-bound ant and transfer it to a different location, it will proceed in the direction where the nest would have been, had we not relocated it (Wehner & Srinivasan, 1981).

In path integration, a creature encodes objects' locations through  $\langle \textit{distance}, \textit{bearing} \rangle$  tuples that specify relations from the creature's current location (and orientation). For instance, a desert ant may represent its nest to be 200 steps (distance) behind it (bearing). As it moves, the ant must update its representation to keep track of the object. Proprioceptive information about the number of steps taken guides these updates, which helps explain how an ant may succeed at navigating without external cues (Wittlinger et al., 2006). Path integration also accounts for the ant's behaviour when picked up: here, proprioceptive information is lacking, the ant cannot update its representation, and hence continues to navigate as if it had happened.

Path integration involves the updating of *egocentric* spatial representations. In such representations, the origin of the represented space is necessarily the representer itself, and all objects' locations are given in relation to this source (Grush, 2000; Klatzky, 1998). A  $\langle \textit{distance}, \textit{bearing} \rangle$  tuple, for instance, necessarily represents the distance and bearing *from the representer's current location (and orientation)*. A creature that represents various objects with such a representation arranges these objects around itself. One object may be *just behind*, another *some distance towards the front*, and yet another *far away to the left*. All representational tokens define relations that originate in the system itself, and they do not represent any relations between (non-self) objects.

As there exists, in the representation, a token that stands for the distance and bearing, these are *explicitly* represented. Something is explicitly represented 'if and only if there actually exists in the functionally relevant place in the system a physically struc-

tured object [...] for which there is a semantics or interpretation, and a provision (a mechanism of some sort) for reading or parsing the formula' (Dennett, 1982, p. 216). The ant's representation fulfils this condition: it contains a *<distance, bearing>* token, and (presumably) there exists a mechanism with which the ant can parse this token to employ it in navigation.

In contrast, the fact that the *<distance, bearing>* token tracks *a relation between the system's location (and orientation)* and some object is only *tacitly* represented. The representation isn't *about* the system itself, it merely *concerns* it (Perry, 1993). While the ant employs the representation to plot a path for *itself* to the nest, it does so without parsing the token to determine that it represents a relation to itself; it automatically employs *all* tokens as representing such relations. The desert ant doesn't represent its own location explicitly, but it knows how to use its egocentric representations to compute how it (and only it) can get to places. While the ant represents a relation, it doesn't represent it *qua* relation (after all, we're missing a *relatum*). The fact that a relation is represented is tacit in the ant's use of the representation.<sup>1</sup>

Thinking about what makes the ant's representation successful or accurate is another way to arrive at the same conclusion. Quite obviously, if the nest weren't where the ant represented it to be, the representation would be inaccurate. But that's not the only way such a representation can fail. If a different ant (at a different location) instantiated the same representation, the representation would likewise fail to successfully guide behaviour. The ant's state – the tacitly represented state implied by how the ant uses the representation – is hence an *unarticulated constituent* of the representation (Perry & Blackburn, 1986; see also Musholt, 2015).

As the representations employed in path integration do not contain a token standing for the system itself, they cannot underwrite an explicit differentiation between self and other. As just mentioned, in such egocentric representations, *all* tokens are used indiscriminately as carrying information concerning relations to the system. These representations necessarily concern the self, but they are not about the self. Such representations therefore cannot give rise to self-representation (even of the minimal kind). They are, strictly speaking, selfless (Lewis, 1979; see also Recanati, 2024).

To find an instance of genuine self-representation, we need to look towards creatures whose more sophisticated behavioural profile relies on non-egocentric representation. In a recent experiment, Tsoar et al. (2011) captured fruit bats in their cave and released them in, or just outside, a crater about eighty kilometres away. The bats that found themselves within the crater, and unable to sense any of their home range's landmarks, engaged in random exploratory flights around the crater. Only

---

<sup>1</sup> What I, following Dennett (1982), call *tacit* representation is sometimes called *implicit* representation in the philosophical and psychological literature (Dienes & Perner, 1999; Musholt, 2013). Additionally, in the psychological literature, 'implicit' and 'explicit' often denote conscious and unconscious or automatic and nonautomatic processes. This is a different distinction.

once they had overcome the crater’s cliffs did they head home. Bats released outside the crater did not engage in random exploratory flights and flew home immediately.

Tsoar et al. (2011) take the experiment to show that bats use *cognitive maps* (see O’Keefe & Burgess, 1996; O’Keefe & Nadel, 1978; Rescorla, 2017) to navigate their environment. We may think of these as representing spatial relations in a manner akin to our phones’ navigation apps, with various tokens standing for the represented objects and the relations between tokens giving their spatial relations.<sup>2</sup> Importantly, such maps allow a system to represent spatial relations between non-self objects, enabling behavioural abilities unattainable otherwise.

With cognitive maps, bats may, for instance, plot novel routes to goal locations even when they aren’t currently given in perception. Once they have determined the distance and bearing to landmarks such as hills and lights emanating from human settlements, bats can use their cognitive maps to compute a direct route home. The experiment’s careful exclusion of various alternative explanations suggests that this is indeed what happened. The crater was outside the bats’ home range so they couldn’t have formed associative links connecting locations along a homeward trajectory. Moreover, because the experimenters transported the bats to the crater, the bats couldn’t have used path integration to keep track of their home. The experimental design also strove to exclude the possibility that bats used magnetic, olfactory, or celestial navigation. The experiment hence suggests that bats represent the spatial relations between locations.

Cognitive maps are a type of non-egocentric spatial representation (Filimon, 2015; Grush, 2000; Klatzky, 1998). In a non-egocentric representation, the origin of the represented space isn’t the representer itself. Instead, the origin might be some (non-self) object (what Grush calls *object-centered reference frames*). Here, the bat would represent the various objects as relations to, for instance, the cave, so the fruit tree is now *in front of the cave*. Alternatively, the represented space may originate at a location without an object (what Grush calls *virtual points of view*). Maps that show the territory from above are one example; each place is shown as a relation to a point that is situated somewhere above the ground.

What is crucial for our purposes isn’t so much the location of the non-self origin of the space, but rather that non-egocentric representations explicitly relate (non-self) objects to one another. One representational token might be about the cave and another one about a distant hill, and the relation between them stands for their spatial relation. Unlike in the egocentric case, these tokens do not tacitly represent – that

---

<sup>2</sup> Cognitive maps do not need to represent distances as straightforwardly as our everyday maps. As research on *predictive maps* argues, distances might instead represent the probabilities associated with getting from one place to another (Cothi et al., 2022; Stachenfeld et al., 2017; Stoewer et al., 2022). According to such a view, token *A* being close to token *B* indicates that the system likely transitions from location *A* to *B*. I believe such maps to be compatible with my proposal, but I cannot argue for this claim here.

is, concern – the location of the representer itself. A non-egocentric representation represents spatial relations as in principle independent of the system’s location. The system’s location does not need to (explicitly or tacitly) figure in the representation, and locomotion hence need not entail a need for accuracy-preserving updates.

To navigate with a non-egocentric representation, you need to know where *you* are. To plot a route home, a bat must first exit the crater and figure out how its own location relates to at least some of the represented objects. It must add to the cognitive map a token that stands for itself. When the bat does so, it instantiates a representation that explicitly attributes properties to itself, distinguishing between its own location and the locations of other objects. Hence, the representation doesn’t merely concern the self but is *about* the self.

However, self-representation requires more than a representation that includes a token specifying the system’s location. If all that is present is such a token, the representation could be about the system, but the system itself might not be aware of that. Imagine some wily experimenter placing a mirror in front of a bat, and the bat, on seeing its own reflection, represents that bat (that is, itself) as located next to a tree. If the bat is ignorant of the fact that it’s seeing itself in the mirror, it won’t be able to use the represented location to, say, plot a route to the tree next to it.

What resurfaces here is a distinction familiar from the literature on *de se* thought, namely between thoughts that happen to be about oneself and thoughts where the agent knows that this is the case. Only in the latter case does the subject think of herself *qua* herself. Consider this seminal example by Perry. He is in the supermarket when he notices a trail of sugar. He forms the belief that someone has a torn sack on their cart. As he follows the trail around the aisles, it eventually dawns on him: he is the one with the torn sack (Perry, 1979). Since we might say ‘Perry believes that he has a torn sack’ to describe Perry’s belief before *and* after his realisation, and since these two states are distinct, it’s become established to distinguish the second case by appending a \* to the relevant pronoun in indirect speech: Perry believes he\* has a torn sack (Castañeda, 1966).

I will introduce a similar bit of jargon to distinguish between self-tokens and self\*-tokens, where, roughly speaking, the former involves a non-egocentric representation with a token that is merely about myself and the latter additionally involves knowledge that this is so. With a mere self-token, a non-egocentric representation has *de re* or *de dicto*<sup>3</sup> content about the self. Such content doesn’t suffice for genuine self-representation, that is, representation with *de se* content. Such a representation requires that the system know which token refers to itself, that is, that it instantiates a self\*-token. Since my interest is in self\*-tokens rather than mere self-tokens, and to keep the terminology as simple as possible, I will use ‘self-token’ to refer to self\*-tokens, except when the distinction becomes relevant again in later sections.

---

<sup>3</sup> The distinction between *de re* and *de dicto* content doesn’t matter for our purposes.



Knowing that some token is about oneself means being disposed to use that token in a specific way when inferring information about current sensorimotor contingencies and when updating the non-egocentric representation based on sensorimotor information. In the case of a spatial representation, such sensorimotor information can be inferred from the vector whose origin is the self\*-token and which ends at the goal location. For instance, when the bat represents itself (self\*-token) to be such-and-such a distance away from the cave (object-token), the bat can infer the sensory input (for instance, visual impression of a dark looming hole) expected given certain motor commands (for instance, moving the wings in a certain fashion for some time). Similarly, if the bat has sensorimotor information regarding, say, a tree given in perception, it can add the tree to its cognitive map by converting the sensorimotor information to a vector which originates at the self\* token and terminates at a token representing the tree.

Sensorimotor representations are a kind of egocentric representation. A sensorimotor representation links motor outputs to sensory inputs, allowing a system to, for instance, calculate the motor outputs required to bring about some desired sensory input. The self doesn't explicitly figure in such representations, but the system's state *does* enter into their accuracy conditions. If we instantiated a given bat's sensorimotor representation in another bat (differently located and/or embodied), the representation would fail to capture the links between motor output and sensory input.

Borrowing a term from Grush (2000; see also Ismael, 2008), I call a non-egocentric representation *coordinated* when a self\*-token is used to establish links to and from sensorimotor representations.<sup>4</sup> Roughly speaking, coordination is the lining up of the non-egocentric representation with the rest of the cognitive system, such that the representation entails, and is entailed by, sensorimotor information. As it is the relations between self\*-tokens and object-tokens that entail, and are entailed by, egocentrically represented information (such as sensorimotor information), coordination necessarily involves a self\*-token. Coordination entails that the system is now disposed to use the non-egocentrically represented information to guide its behaviour and to use information from the senses to update its non-egocentric representation.

A coordinated non-egocentric representation fulfils the two conditions on self-representation since it represents a system's properties explicitly and does so in a way that directly links the represented properties to behaviour and sensation. For instance, a bat that non-egocentrically represents itself to be close to a food source realises a representation that contains, as a minimum, tokens for itself and the food source. This bat explicitly differentiates between its location and another object's location. Moreover, if the representation is coordinated, then the bat is disposed to use relations between the self\*-token and object-tokens to infer sensorimotor information that guides its behaviour. Finally, the bat is disposed to use information from its senses to update the

---

<sup>4</sup> My use of the term 'coordination' should not be conflated with Grush's (2000). What I call coordination is only one variant of Grush's more general kind.

non-egocentric representation.

So far, I've only given a sketch of the argument. The rest of the paper fills in the gaps. Over the next few sections, I introduce instances of non-egocentric representation of temporal, bodily, and cognitive states. Following this, I argue, in quite some detail, why such representations are all genuine (if minimal) self-representations.

## Non-egocentrically representing temporal states

Various authors have already discussed a distinction akin to the one I draw between egocentric and non-egocentric representation in the domain of temporal cognition (Grush, 2008; Kort et al., 2005; McCormack, 2015). Moreover, many authors think that agents need to – or at least often do – self-represent temporal properties (Campbell, 1999; Goldie, 2012; Lamarque, 2004; Menary, 2008; Peacocke, 2014), making it important that I show that non-egocentric representations can account for these. I'll proceed in two steps: First, I look at egocentric representations of temporal properties and show how these fail to underwrite certain more complex forms of behaviour. I then look at non-egocentric representations, discuss what behaviours require them, and tease some links between such representations and self-representation.

When food becomes scarce, slime mould amoebae merge into plasmodia that move around to forage for nutrients. Saigusa et al. (2008) exposed these plasmodia to multiple pulses of dry conditions, capitalising on the well-known fact that they slow down in dry environments. They found that after just three exposures, slime moulds periodically slow down even if the dry condition stimulus is removed, which indicates that they have started to *anticipate* the dry conditions.

To behave anticipatorily, slime moulds need to keep track of time, which they likely do with the help of regular internal oscillations (Saigusa et al., 2008) similar to those employed by various other creatures (Gallistel, 1989). Slime moulds can use these oscillations to encode a *timespan* variable tracking the temporal distance to the next dry condition. Using such *interval timing* (Buhsu & Meck, 2005), a slime mould may then slow down its forward movement when an impending dry condition is indicated by *timespan* approaching zero.

Significant parallels between path integration and interval timing suggest that both use egocentric representations. Where *timespan* specifies the temporal distance of some event *from the present*, *distance, bearing* specifies the spatial distance and bearing *from the current location (and orientation)*. Moreover, in both cases, the representations merely concern (and aren't about) relations *from the system's actual state*. The system's actual state doesn't explicitly figure in the representations but enters as an unarticulated constituent that helps determine the representations' accuracy conditions. Moreover, because such representations *necessarily* concern relations to the system's actual state, egocentric temporal representations – just as ego-

centric spatial representations – need to be updated with each change in the system’s state. While path integration is employed to update spatial representation, egocentric temporal representations are kept accurate using internal oscillations that decrement (*timespan*).

There are restrictions on what egocentric temporal representations can represent since they must always represent temporal properties with a single token that concerns a relationship to the present. For one, they ‘are not carrying contents to the effect that this or that thing is happening at this or that time’ (Grush, 2008, p. 156). They cannot explicitly represent relations between moments in time and thus cannot represent that an event occurs at *some* moment in time. Consequently, egocentric representations of temporal properties cannot represent an event as happening before or after another (except tacitly, when one of the events is in the present). McCormack (2015) calls this a conception of *time as duration*, aptly capturing that the explicitly represented content is a duration from the present to some other moment in time.

Both Grush and McCormack contrast this form of temporal cognition with ‘a more sophisticated kind’ (Grush, 2008, p. 156) that does succeed at explicitly representing relations between events. A study by Arcediano et al. (2003) illustrates this more sophisticated form of temporal cognition (see also Barnett et al., 1997; Taylor et al., 2014). They first exposed rats to two neutral stimuli (two sounds in this case) following one another with a five-second gap ( $S_2 \rightarrow S_1$ ). In a second phase, the same rats were exposed to an aversive stimulus *US* (a footshock) paired with  $S_1$  ( $US \rightarrow S_1$ ). In the test phase, the rats were exposed to  $S_2$ . The rats displayed decreased use of a lick bottle following this exposure, demonstrating a fear response.

Note how the rats seemed to anticipate the footshock on exposure to  $S_2$  even though they had not been conditioned on the relevant sequence ( $S_2 \rightarrow US$ ). The study’s authors assert that the rats’ capacity to anticipate without being exposed to the relevant sequence of events shows that rats can integrate the two learnt temporal sequences to arrive at a sequence  $S_2 \rightarrow US \rightarrow S_1$ , leading the rats to expect *US* on being exposed to  $S_2$ . Arcediano et al. (2003) think this integration requires a kind of cognitive map already suggested by Tolman (1948), a temporal map.

Here’s why we might think that egocentric representations cannot explain this behaviour and why a non-egocentric representation, such as a temporal map, is needed. A creature could use an egocentric representation of temporal properties to measure the time between the events taking place in the first two phases of the experiment. In the first phase, it could instantiate a representation  $R1_{ego}$  to measure the duration between  $S_2$  and  $S_1$ . In the second phase, another representation  $R2_{ego}$  would measure the duration between *US* and  $S_1$ . Both representations only tacitly represent the fact that they are relations *between* certain events. More specifically, the fact that these representations are about temporal relations is tacit in their activation at  $S_2$  and *US*, respectively, and their subsequent employment in behaviour anticipatory of  $S_1$ . That’s why, even though  $R1_{ego}$  and  $R2_{ego}$  carry explicit content about a relation to  $S_1$ , the rats

cannot, say, subtract the value in  $R_{2_{ego}}$  from  $R_{1_{ego}}$  to arrive at the temporal distance between  $S_2$  and  $US$ . Or, more precisely, while they *could* subtract one value from the other, they would not thereby have any indication of how to use the resulting representation; they would not, unlike in the case of  $R_{1_{ego}}$  and  $R_{2_{ego}}$  have formed a non-representational association between a certain event (here,  $S_2$ ) and activation of the token.<sup>5</sup>

To integrate the two temporal sequences, rats must represent them non-egocentrically. When their representations  $R_{1_{non-ego}}$  (about  $S_2 \rightarrow S_1$ ) and  $R_{2_{non-ego}}$  (about  $US \rightarrow S_1$ ) explicitly represent relations between events, they may integrate representations that share an element (Gür et al., 2018; Molet et al., 2012; Taylor et al., 2014). In the experiment, rats first learn that  $S_2$  precedes  $S_1$ . Then, in the second phase, they learn that  $US$  precedes  $S_1$ . Because  $S_2$  doesn't occur between  $US$  and  $S_1$ , they seem to infer that  $S_2$  occurs before  $US$ , arriving at  $S_2 \rightarrow US \rightarrow S_1$ .

In a non-egocentric representation, time functions as a framework (McCormack, 2015) that organises the temporal relations between events. These relations do not tacitly refer to the system's actual state (the present). The absence of such reference is evidenced by the rats' representation of  $S_2 \rightarrow US \rightarrow S_1$  remaining accurate as time passes even if it isn't updated. In short, temporal relations are represented in a way that is in principle independent of the system's actual state.

Since non-egocentric temporal representations do not necessarily concern the system's own state, coordination becomes an issue. Just as a creature may sometimes be ignorant about where it is, it may also be ignorant about *when* it is. The rats in the experiment illustrate this: They may store a representation of  $S_2 \rightarrow US \rightarrow S_1$  in memory without knowing how those events temporally relate to the present. It may only be at, say, the occurrence of  $S_2$  that the rat is able to infer that *it* is now located at the same moment in time as  $S_2$ . It can then infer sensorimotor information from the non-egocentric representation and behave in a manner anticipatory of  $US$ , suppressing the use of the lick bottle.

Note how, just as in the case of spatial cognition, coordinated non-egocentric representation of one's temporal properties seems to be linked to self-representation. A creature with such a non-egocentric representation will explicitly differentiate between its temporal properties and those of events in the past and future. Such a representation is about the system instantiating the representation. Moreover, through coordination, the system uses a certain token in a special way, namely as a self\*-token that specifies the system's actual temporal state. The system is disposed to infer sensorimotor information from relations between the self\*-token and other

---

<sup>5</sup> We might attempt a different explanation: As the rats associate  $S_1$  with the footshock ( $US$ ) in the second phase of the experiment,  $S_1$  takes on a negative valence. Later on, when the rats are exposed to  $S_2$  in the test phase, it's the expectation of  $S_1$  rather than  $US$  that causes the fear response. Arcediano et al. (2003) are aware of this possibility and implemented an alternative version of the experiment to exclude it. I refer the interested reader to their paper for more information.

tokens and to update the representation based on sensorimotor information. If this is right, a creature instantiating non-egocentric temporal representations is a genuine self-representer.

## Non-egocentrically representing bodily states

Many animals represent their own bodies. These body representations develop early in ontogenesis (Meltzoff et al., 2019; Slaughter & Brownell, 2011) and are said to prepare the ground for later representations of the body-external environment (Nguyen et al., 2020; Stoytchev, 2009). Some of these representations are non-egocentric and suggest a link to minimal forms of self-representation. Unlike in previous sections, and in the interest of brevity, I will directly turn to non-egocentric representations without contrasting them with their egocentric brethren.

Human infants begin to successfully reach for objects at around three to four months of age. At first, infants reach jerkily, inaccurately, and with movements that do not describe a straight line (Bruner & Koslowski, 1972; Hofsten, 1982; Thelen et al., 1996). Such movements expend more energy than is strictly speaking required to reach the target state (Neil E. Berthier & Keen, 2006; Konczak et al., 1995; Rolf & Steill, 2012). Starting at around five months of age, infants anticipatorily adjust their hands to the shape of the object to be grasped at the end of the reach (Hofsten & Fazel-Zandy, 1984; Wentworth et al., 2000; Witherington, 2005), and by around seven months of age, infants can reach for objects efficiently and effectively (Hofsten, 1991; Thelen et al., 1996). They are now, for instance, able to drink from an open cup (Hofsten & Fazel-Zandy, 1984). At about two years, smoothness of motion approaches adult abilities (Neil E. Berthier & Keen, 2006).

Originally, researchers explained the early (jerky, inaccurate, zig-zag) movements with the supposed vision-guided nature of infant reaching (for an overview of that literature, see Neil E. Berthier (1996) and Corbetta et al. (2018)). It was believed that infants would pause their reaching behaviour at various points along the trajectory to check and correct their motions by gazing back and forth between the hand and the intended target. Only later would infants develop the embodied sense of their own body that enables smooth bodily movements.

In light of several findings, this view has been largely abandoned. Hofsten & Lindhagen (1979) observed that infants do not shift their gaze back and forth between target and hand position but rather fixate on the object. Even more convincingly, Clifton et al. (1993) showed that infants can successfully reach for objects even when they cannot see their own hands and arms. According to these results, when infants start engaging in reaching behaviour, they must *already* have an embodied sense of their hand location (Corbetta et al., 2014).

Converging evidence from developmental psychology and robotics (Baranes &

Oudeyer, 2013; Castellini et al., 2007; Desmurget & Grafton, 2000; Rolf & Steill, 2012; Schillaci et al., 2016; Takemura et al., 2018) indicates that we can explain infants' developmental trajectories and emergent abilities with internal body models, *forward models* and *inverse models*. Forward models (Grush, 2004; Körding & Wolpert, 2006; Wolpert & Ghahramani, 2000) allow systems to compute the bodily (and sensory) state to which a motor command will likely give rise. As forward models represent the bodily dynamics internally, a system can use them to estimate its state even when sensory input is noisy. Furthermore, the predictions of forward models are realised in the brain and are accessible long before proprioceptive signals have completed their journey back from the sensory surfaces. With a forward model, an infant could, say, predict that the motor commands she just issued will cause her to fail to reach a toy in front of her. Importantly, the infant could do so in the absence of – or before – visual sensory input that confirms the failure of the reaching movement. Inverse models are, as the name implies, the inverse of forward models. They allow computing the motor commands required to reach a target state and are hence implicated in goal-directed behaviour. An infant could use her inverse model to calculate the motor commands required to reach the toy without visually aligning her hand with the target.

Forward and inverse models need information about the system's current bodily state to compute their outputs. For instance, depending on whether her arm is flat against the side or stretched out in front of her, an infant's inverse model should issue different predictions about the motor commands necessary to reach a toy. The forward model, too, will predict different bodily states depending on the infant's current posture. Hence, the forward model relates possible bodily states in terms that may be inferred from motor commands, whereas the inverse model relates possible bodily states in terms that allow inferring the motor commands required to transition between the states.<sup>6</sup>

Internal body models are non-egocentric representations. Not only do internal body models explicitly differentiate between various possible states, but the relations that these bodily states are represented to bear on one another are in principle independent of the system's actual posture. When such a representation contains a token giving the system's actual state, it is hereby *about* the system's actual state.

To employ such a model, a system needs to use one of the body state tokens in a special way, namely as specifying its own actual state. Or, to put it in the words of Coslett et al. (2008), 'you are here' information is 'required in order to accurately reach toward an object' (p. 117). In other words, an internal body model needs to be coordinated for the system to use it. Only then is the system disposed to infer sensorimotor

---

<sup>6</sup> Forward and inverse models only explicitly differentiate between possible bodily states if *articulated* (see Grush, 2004). Such articulation is likely the case. Grush references research showing correlations between certain groups of neurons and, for instance, elbow angle. Wolpert and Ghahramani note that the human muscular system comprises over 600 muscles and that the resulting complexity 'clearly prohibits a simple look-up table' (Wolpert & Ghahramani, 2000, p. 1212).

information from relations between the self\*-token and tokens that represent merely possible bodily states. And only then is it disposed to infer possible bodily states from sensorimotor information.

Coordinated non-egocentric representation of bodily states suggests a form of self-representation. An inverse model non-egocentrically represents bodily properties by representing the system as exemplifying a certain posture that is explicitly differentiated from other postures the system could exemplify if it were to issue certain motor commands. It is about the system's actual state (and doesn't merely concern it). Reminiscent of certain non-egocentric temporal representations, the *other* from which the system is explicitly differentiated here is the system's own merely possible state.<sup>7</sup>

## Non-egocentrically representing cognitive states

We routinely represent ourselves to have beliefs, desires, and other mental states. Moreover, some authors even think that representing one's mental states is necessary for self-representation (see, for instance, Musholt, 2013, 2015). It's hence important for my account to explain how non-egocentric representations can underwrite self-ascriptions of mental and cognitive properties. This section shows that there aren't any principled difficulties in extending my account to such properties. Additionally, we'll see that such representations are even instantiated by certain non-human animals.

Corvids, and in particular scrub-jays, have recently been the subjects of a slew of experiments that highlight their sophisticated forms of social cognition (Baciadonna et al., 2022; for instance, Bugnyar et al., 2016; Kort et al., 2005, 2005). Many of these studies cleverly exploit scrub-jays' predilection for caching food and pilfering conspecifics' caches. One such experiment by Clayton et al. (2007; see also Nathan J. Emery & Clayton, 2008) studied how cachers adapt their behaviour when a conspecific is present during cache recovery. At a first caching event, a conspecific (observer A) who looks on through a transparent partition observes the cacher. On the cacher's side, two trays – visible to both birds – constitute potential caching locations, though one of the trays (tray B) is inaccessible due to a transparent obstacle. The obstacle forces the cacher to hide the food in tray A. At the second caching event, a different

---

<sup>7</sup> The fact that mirror neurons may cause babies' (and adults') body representations to also activate when they watch another person reach or grasp complicates the picture (Lepage & Théoret, 2007; Marshall & Meltzoff, 2015). Represented bodily states are then only sometimes the infant's own, and this has been argued to underlie their ability to imitate others. The idea is, roughly speaking, that infants automatically mimic others because they (sometimes) fail to differentiate between the targets of perceived actions and their own target states. However, note that the infant still needs to use some token as a self\*-token and that this token cannot be what is confused with the perceived state. Since the child would perceive herself as already embodying the target state, observing the other would not entrain motor behaviour if she confused the other's state with her own.

conspecific (observer B) observes the cacher. This time, tray A is inaccessible, and the cacher hides food in tray B.

At recovery, the cacher behaves differently if observer A rather than B is present. For instance, when observer A is present, cachers tend to re-cache the food from tray A under tray B, often moving it back and forth several times first. They will not recover the food from tray B. Presumably, this behaviour aims at better protecting the food initially stored under tray A (that the observer saw being cached) while not giving away the location of the food under tray B (that the observer did not see being cached). When observer B is present, the behaviour is analogous but targets the other tray.<sup>8</sup>

The study's authors believe the experiment reveals that corvids employ representations of conspecifics' mental states, though others think that such conclusions are premature (Derek C. Penn et al., 2008; Derek C. Penn & Povinelli, 2007). The experiment fails to show, they say, that corvids exhibit 'a sensitivity to what others have and have not seen' (Clayton et al., 2007, p. 519) or that they engage in 'knowledge attribution' (ibid.). The caching bird could, for instance, simply represent which observer had been present at which caching event. At recovery, the cacher could then re-cache food it hid while this particular observer had been present. We can then explain scrub-jay behaviour by positing that corvids represent spatial and temporal properties rather than conspecifics' cognitive states.

However, the fact that only birds that have pilfered others' caches engage in cache protection suggests the need for an alternative explanation (N. J. Emery & Clayton, 2001). Corvids seem to infer others' behavioural patterns from their own, which requires that they distinguish between their own informational states and others' informational states. Only then can the cacher differentiate between where it thinks the cache is located and where the conspecific thinks it is located. Moreover, the cacher must conceive of the informational states attributed to itself and the conspecific as being of the same kind. Only then can the scrub-jay be sensitive to the fact that these states determine behaviour in the same sort of way. In short, scrub-jays need to be able to represent that conspecifics may be in informational states of the same *kind* but differing in *value* from their own.

The spatial representations covered in previous sections were always tacitly about how the *system itself* estimated the environment to be. As we just saw, they are insufficient to explain corvid caching behaviour. Butterfill & Apperly (2013; see also Apperly & Butterfill, 2009) argue that scrub-jays may distinguish between their own and others' representations of spatial properties without representing mental states. Instead, they represent *registrations*, which are representations of others' relations to certain objects: 'an individual registers an object at a location if and only if she most recently encountered it at that location' (Butterfill & Apperly, 2013, p. 617). Like beliefs, registrations motivate behaviour and may be incorrect (when the object is no longer

---

<sup>8</sup> Clayton et al. (2007) went to considerable lengths to exclude the possibility that cache protection behaviour is cued by the observer's behaviour. More details can be found in their paper.



where it had been registered). A scrub-jay may use representations of registrations to predict others' behaviour across a wide (though limited<sup>9</sup>) range of situations. Seeing that some conspecific has registered food at a certain cache, a scrub-jay could thus infer that the conspecific will attempt to pilfer it.

Scrub-jays, according to this proposal, represent registrations by encoding several  $\langle individual, location, object \rangle$  tuples. At cache recovery, a cacher first selects the registration (or registrations) whose *individual* corresponds to the conspecific currently present and whose *location* is the current location. The cacher then attempts to make it so that the *location* and *object* elements of the registration, which indicate the conspecific's registration of the *object* at the *location*, do not correspond to its own representation of cache location. For instance, if a scrub-jay represents the food to be located under tray A while representing the conspecific to have registered the food under that same tray, it will remove the food from the cache. The food is subsequently represented as being in a location that is different from where the conspecific registered it.

We still need to shed light on how a scrub-jay compares a conspecific's registrations with its own representation of cache locations. After all, and as I've mentioned above, for the cacher to infer the observer's behaviour from its own, the informational states involved in the two cases need to be comparable. Registrations aren't of the same kind as the scrub-jay's own representations of spatial properties and therefore cannot be directly compared to these. The scrub-jay must therefore first convert its representation of spatial properties into  $\langle individual, location, object \rangle$  tuples, where the *individual* is the scrub-jay itself. It must create a representational token that refers to itself and with which it attributes to itself certain registrations. These registrations can then be compared to those of conspecifics. How this is done doesn't matter for our purposes; what matters is that translation is required and that it involves self-attributing registrations.

The resulting representation should be familiar from earlier sections: it includes tokens standing for objects (individuals) and their properties (registrations), and to employ the representation, the system must use a token as specifying its own state. This structure evidences the hallmarks of coordinated non-egocentric representation: object properties are represented as in principle independent of the system's own state, and use of the representation requires coordination.

There are intriguing links to self-representation, much like with non-egocentric representations of other types of property. A scrub-jay explicitly differentiates between where *it* has encountered objects and where conspecifics have encountered them. The representational token that specifies its own registrations is *about* itself. Relations between its own registration and those of others imply certain sensorimotor information as can be seen in scrub-jays' flexible caching behaviour. Such

---

<sup>9</sup> Registrations do not exemplify beliefs' full functional profile (Butterfill & Apperly, 2013). These details are not relevant to my argument.

non-egocentric representation therefore seem to imply a form of self-representation – scrub-jays seem to represent themselves (and others) as exemplifying cognitive states.

Note that we might still doubt whether the representation of one's registrations entails that scrub-jays self-represent properly *mental* properties (Burge, 2018). We might instead say that these are self-representations of merely *cognitive* states. If this is so, authors such as Musholt (2013), who believe that self-representation requires representing one's mental states, would claim that corvids do not self-represent. I believe that this is false: any coordinated non-egocentric representation is a self-representation. I now turn to giving substance to this claim.

## Minimal self-representation

I've argued that coordinated non-egocentric representations are used by various creatures to represent a diverse set of properties, and I've hinted at connections between this kind of representation and *self*-representation. I now move beyond hints and show why coordinated non-egocentric representation is minimal self-representation. Doing so requires showing how non-egocentric representations fulfil the two conditions on self-representation. First, they must explicitly attribute properties to the representer itself. Second, the representer must know that the representation is about itself, where this manifests in a direct link to behaviour and sensation.

As mentioned before, representations that explicitly distinguish between self and other are representations that are *about* – rather than merely *concerning* – the self. A bat that represents its location in a cognitive map instantiates a representation about itself, whereas the ant's representation of the relation to various points of interest only concerns it. In the account I propose, this is captured by non-egocentric and egocentric representations, respectively.

Egocentric representations necessarily represent the self tacitly. They concern the self and are not about the self. An ant that represents the distances and bearings to its nest and some food source instantiates representational tokens whose  $\langle \textit{distance}, \textit{bearing} \rangle$  tuples necessarily concern relations to the ant's actual state (in particular, the ant's location and orientation). Such a representation doesn't contain tokens that stand for the ant and, hence, isn't about the ant. Egocentric representations of other properties are no different. As we've seen, slime moulds encode  $\langle \textit{timespan} \rangle$  variables, which necessarily concern the temporal distance of some event to the slime mould's actual state (here, the temporal present). However, such a representation doesn't contain a token that stands for the slime mould and, hence, the self only figures as an unarticulated constituent. The format of egocentric representation is, as mentioned before, selfless but necessarily self-concerned (Lewis, 1979; Recanati, 2024).

In contrast, non-egocentric representations do not involve the self as an unarticulated constituent, resulting in the representation of objects and their properties as in principle independent of the system's own state. A fruit bat may fly from the crater to the cave without updating the tokens representing these locations. A scrub-jay may update its representation of cache locations (or, more precisely, its representation of its own registrations) without needing to update its representation of others' registrations.

To be *about* the system, and fulfil the first condition on self-representation, a non-egocentric representation must represent the system's own state. It does with an object-token that stands for the system, which I call (mere) self-token. With such a token, the fruit bat may represent its location and the scrub-jay its registrations. Here, the self is an *articulated* constituent, and relations between the self-token and other tokens issue in explicit self-other differentiation. Note that in some cases, such as internal body models, the self-token relates to object-tokens that refer to merely possible states the system might exemplify. This complication doesn't change anything of substance – in either case, we now have a representation *about* the (actual) self, which fulfils the first condition on self-representation.

However, with only a (mere) self-token, a system fails to know that it\* is the represented object, which precludes genuine self-representation. A fruit bat may have a maximally detailed cognitive map that includes a token specifying its own location, but if it doesn't know that this token stands for its\* location, it cannot use the representation for locomotion, and it doesn't self-represent. Similarly, a rat representing a sequence of events may instantiate a token that stands for the present, but if it doesn't know what token is about its\* temporal state, it cannot use the representation to behave anticipatorily, and it doesn't self-represent.

Self-representation requires that the non-egocentric representation is directly linked to sensation and behaviour. The bat needs to know where it\* is and the rat needs to know when it\* is. This is akin to the examples discussed in the literature on self-locating beliefs. Perry needs to realise that he\* is the one with the torn sack of sugar before he knows to adjust the sack on his cart to stop it from spilling (1979), and Mach needs to realise that he\* is the dishevelled pedagogue to, say, reach for the comb to tame his unkempt hair (1890).<sup>10</sup>

---

<sup>10</sup> Note that a creature may, in certain situations, employ even an uncoordinated non-egocentric representation. Imagine that fruit bats can tell one another the location of food sources. A bat could then tell a conspecific that some fruit tree is closer to the cave than another. To do so, the bat doesn't need to know its location. However, while the contribution of the non-egocentric representation may be important, even necessary, it isn't sufficient. The bat cannot infer the relevant sensorimotor contingencies from the uncoordinated non-egocentric representation alone but needs additional sensorimotor information about the communicative situation. For instance, if it knows that it can expect to be treated favourably in the future if it helps the conspecific, it could then use the non-egocentric representation to figure out how to help. For our purposes, it suffices that coordination *is* required in many important cases.

According to my proposal, a non-egocentric representation becomes a self-representation when it is *coordinated*. Coordination establishes systematic non-representational – or architectural (Ismael, 2008) – relations between sensorimotor representation and non-egocentric representation so that the creature is disposed to infer sensorimotor information (in an egocentric format) from non-egocentric representation and vice versa.

Such coordination necessarily creates a special sort of token, the self\*-token. The self\*-token is the representational token with which a system represents itself and does so in a way that directly links it to behaviour and sensation. A self\*-token is, as mentioned before, distinct from a (mere) self-token. A self\*-token isn't individuated by the object for which it stands but rather by the special way the rest of the cognitive system *employs* it. Think of spreadsheet software, which presents a view of cells arranged in rows and columns. When you choose a cell, it becomes highlighted. If you then type a command, it only impacts that particular cell. The highlighting indicates how the interface is linked to the data; it indicates on what piece of data the user will operate. The highlighting doesn't add any information to the spreadsheet but merely establishes a link between the spreadsheet view and the underlying data. Like the highlighting in such software, a self\*-token doesn't add information to non-egocentric representations but indicates which (pre-existing) token is linked to the rest of the system in the special way that manifests the system's know-how that this token stands for itself.

A system knows that the self\*-token is about its\* state when it employs the token in inferences to and from sensorimotor information. I first look at how inferences from non-egocentric representation to sensorimotor information involve the self\*-token. Sensorimotor information is, as mentioned, represented egocentrically, which means it is necessarily self-concerning. Non-egocentric representation, on the other hand, represents objects as in principle independent of the system's state. Hence, for a non-egocentric representation to imply sensorimotor information, one of its tokens must be *used as* denoting the system's state. If a fruit bat is to infer the motor commands needed to get home, it must instantiate, in addition to the object token about the cave, a self\*-token that indicates its\* own position. A rat may only use its non-egocentric representation to predict future events when it uses one of the represented moments in time as its\* temporal state. And it's only when an infant uses one of the many bodily states represented by an internal model as specifying her\* actual bodily state that she may employ an inverse model to infer the motor commands required to reach some goal state.

Since a creature cannot use her non-egocentric representation without a self\*-token, reference to the self\*-token is necessary to explain how the representation issues in the behaviour. This mirrors the observation, first noted by Castañeda (1966) and Perry (1979), but later also discussed by many others (for instance Babb, 2016; José Luis Bermúdez, 2017; Lewis, 1979; Stalnaker, 1981) that certain indexicals are essential

to explanations of actions.<sup>11</sup> It is impossible, they say, to explain why Perry adjusted the torn sack in his cart without saying that this is so because he realised that he\* is the one with the torn sack. When we replace 'he\*' with (non-indexical) descriptions – maybe 'the only bearded man in the store' – we can no longer explain the change in behaviour (if we don't attribute to Perry the knowledge that he\* is that bearded man). No matter the description, Perry must additionally realise that he\* is the person described. Similarly, to explain why a scrub-jay re-cached the food under a certain tray, we must refer to a self\*-token. No matter what non-egocentrically represented information we give the bird (in the form of object-tokens, even including mere self-tokens), it cannot infer sensorimotor information. Non-egocentric representations represent objects as in principle independent of the system's state, and sensorimotor information cannot be inferred without a self\*-token. Only when the representation is coordinated, and a token is specified as self\*-token, does the cognitive system know the origin of the non-egocentrically represented relation that implies sensorimotor information.

Links in the opposite direction are also crucial for self-representation. These links allow the system to update its non-egocentric representation with information encoded in an egocentric format. A rat with sensorimotor information about some future event can only encode that information in its non-egocentric representation if it knows which representational token is about its\* temporal state. A non-egocentric token's value can only be inferred based on sensorimotor information *and* the self\*-token. Only then can the rat establish a correspondence between the sensorimotor information and a *relation* from its\* temporal state (that is, the self\*-token) to the event. The same applies to the other creatures we have looked at: on having gained sensorimotor information about the motor commands required to, say, lift her arm, an infant may want to use this information to enrich her internal model. Since the sensorimotor information corresponds to a relation between her\* actual state and some goal state, a self\*-token is necessary to update the non-egocentric representation.

I have argued that non-egocentric representations explicitly differentiate between properties attributed to the self and those attributed to others and that coordination establishes links between such a representation and sensorimotor information. These links entail, first, that a creature with a coordinated non-egocentric representation is disposed to take the represented information as directly relevant to its behaviour. Second, the links entail that the creature is disposed to update the representation based on information supplied by the senses.

Any creature representing any property with a coordinated non-egocentric representation is a self-representer. Minimal self-representation does not depend on the

---

<sup>11</sup> Cappelen & Dever (2014) oppose the prevailing opinion and argue that 'there is no such thing as essential indexicality, irreducibly de se attitudes, or self-locating attitudes' (p. 3). Millikan (1990) argues that the relevant mental tokens *are* essential, but not indexical (see also Prosser, 2015). According to her, these tokens are distinguished by their psychological role rather than their semantics. I am sympathetic to this latter view, but space constraints prevent me from discussing this topic in more detail.

representation of specific kinds of property, and as the examples throughout this paper have shown, minimal self-representation is more common than we might have thought. Scrub-jays, fruit bats, young infants, and rats all self-represent some of their properties.

## How to find out who you are

While I've covered why coordination – and the self\*-token – is essential for establishing direct links to behaviour and sensation, we're still missing a crucial piece of the puzzle: how can a system determine which token to use as a self\*-token, that is, how can it coordinate its non-egocentric representation with sensorimotor information? The present section covers the two broad ways a creature may gain information about itself: first, by making the unarticulated subject component of its sensorimotor representations explicit and, second, by identifying itself with one of the non-egocentrically represented objects. I will show how self-attributions of the first kind are immune to certain errors through misidentification, which many philosophers see as essential for self-representation (Jose Luis Bermúdez, 1998; Evans, 1982; Musholt, 2013; Shoemaker, 1968).

The most important (though not only) way a system can learn about its\* properties is by establishing correspondences – making identity judgements – between sensorimotor information and *relations* in the non-egocentric representation. For instance, an infant may possess sensorimotor information linking various motor commands to bodily movements. She can then use this information to search her internal body model for a self\*-token whose relations to various object-tokens imply the sensorimotor information she possesses. A system can infer sensorimotor information from relations between a self\*-token and object-tokens, and by doing the reverse, it can infer a self\*-token from sensorimotor relations.

A creature can establish coordination since sensorimotor information necessarily concerns the self. When a creature matches sensorimotor information to a relation in the non-egocentric representation, the object-token at the origin of that relation articulates the sensorimotor representation's unarticulated subject component. Since the unarticulated component of the sensorimotor representation is necessarily the self, the object-token that is its articulation also necessarily refers to the self. The system can therefore directly, and without the possibility of error, designate that token as a self\*-token.

Since such self\*-tokens cannot attribute properties to the wrong object, they are immune to errors through misidentification (IEM) in the same way as certain uses of the indexical 'I' (Evans, 1982; Perry, 2010; Shoemaker, 1968; Wittgenstein, 2007). These uses of 'I' – what Wittgenstein (2007) called uses of 'I' as subject – make it impossible to ascribe the property to the wrong person. When thinking 'I see a tree' (or 'I have

a headache'), it makes no sense to wonder, 'Someone is seeing a tree, but is it *I* (or 'Someone is having a headache, but is it *I*)? I might be mistaken about the content of the visual experience, but the fact that *I* am having the experience is beyond doubt. The matter is no different with self\*-tokens inferred from sensorimotor information: while they may misrepresent a system's properties, they necessarily refer to the system itself. Coordinated non-egocentric representation 'allows for the possibility that it misrepresents the property that is being ascribed, while it cannot misrepresent the subject purportedly possessing that property' (Musholt, 2013, sec 2.3). A bat might infer its self\*-token from inaccurate sensorimotor information, leading to a misrepresentation of its properties in the non-egocentric representation. However, since sensorimotor information is necessarily self-concerned, the self\*-token cannot but refer to the system itself.

My explanation of why certain self-ascriptions are IEM is similar, in certain ways, to the account advanced by Recanati (2009, 2012a, 2024). Like me, Recanati argues that the status of certain self-ascriptions as IEM derives from the content of the grounds for those self-ascriptions. It's because the contents are of a certain kind that a system need not identify with some object to self-ascribe a property, hence making mistakes through misidentification impossible. However, Recanati further argues – and here we part ways – that self-ascriptions are IEM when grounded in experiences with a certain kind of *mode* (see Searle, 1983). The mode of an experience phenomenally distinguishes experiences of different kinds without being part of the experience's content. For instance, proprioception's functional role is to provide information about one's body, and because of this functional role, such experiences feel a certain way. Importantly, such experiences are selfless so that when I experience my legs being crossed, I just experience *crossed legs*. While the self doesn't figure in the content of the experience, an agent can identify the kind of experience she is having by its mode, and given that experiences of this mode necessarily provide information about the self, '[t]he person in question is, as it were, pre-identified, being determined by the mode of the experiential state' (Recanati, 2024, p. 9). Hence, I cannot misattribute the property when I move from a proprioceptive experience of crossed legs to attributing to myself the property of having crossed legs.

This approach presents a difficulty: how can we explain why experiences of the *external mode* – experiences, such as perceptual ones, whose functional role isn't to provide information about the system itself – can also ground self-ascriptions that are IEM? Recanati argues that this is so because these experiences are also 'bound to be about the subject of experience' (2024, p. 10). When an agent sees the Eiffel Tower, then the Eiffel Tower is necessarily present *in the subject's environment*. Hence, here too, there cannot be any IEM when the agent forms the belief that *I am standing in front of the Eiffel Tower*. However, note that the mode of the experience no longer does any work: no matter whether the experience is of the internal or external mode, it can be grounds for a self-ascription that is IEM.

In my view, self-ascriptions are IEM when their grounds are representations of a specific kind of format, namely egocentric representations, and this is why experiences of the internal and external mode can provide the relevant grounds. Note, first, that the crucial factor isn't that egocentric representations are selfless. After all, when I look into the mirror, I see *myself*, but this perceptual experience can still ground the IEM self-ascription that *I stand in front of a mirror*. What's crucial for IEM is that the relevant representations (or, for Recanati, experiences) are self-concerning. These representations are self-concerning because they are in an egocentric format, and it's this representational format, rather than the mode of the experience, that underwrites IEM (viz. Ismael, 2012). Experiences of the internal and external mode can ground IEM self-ascriptions because these types of experience are realised by (self-concerning) egocentric representations, and the resulting self-ascriptions are IEM when they involve the articulation of the egocentric representations' unarticulated subject components.<sup>12</sup>

The proposed account of self-representation also explains self-ascriptions that aren't immune to errors through misidentification. Whether a non-egocentrically represented property is IEM depends on the grounds for the self-ascription (Evans, 1982). If, after an accident, I (visually) perceive a broken arm and form the belief that my\* arm is broken, this belief isn't IEM. The arm I'm seeing may not be my arm and, hence, the person with the broken arm may fail to be me. Here, my judgement that I have a broken arm is ultimately based on the articulated (explicit) *object* component of a sensorimotor representation. My seeing of the broken arm might instantiate sensorimotor information to the effect that if I moved my head this or that way, I would see it from this or that angle. I could then use this information to add an object-token to

---

<sup>12</sup> We might also criticise Recanati for failing to link his account of IEM with the *mental files* account he espouses (Recanati, 2012b; see also Peacocke, 2014). While the differences between my account and the mental file account are too wide-ranging and subtle to discuss here, I want to take this opportunity to contrast the two in broad strokes. According to the mental files account, agents instantiate mental files which contain information about various objects. Mental files are individuated by *epistemically rewarding (ER) relations* through which we gain information about the object. For instance, on seeing a tree, I might token an object file and store in it information such as *is a tree, is tall, is green*, and so forth. The self-file, in this view, is a standard file that stores information where 'the relevant ER relation is the identity relation' (Recanati, 2012b, p. 68). What this account disregards – or at least fails to foreground – is that self-representations aren't just stores of information but are directly linked, in the manner I discuss, to behaviour and sensation. It is interesting, given this context, that Recanati's account of IEM doesn't refer to his account of self-files, a fact commented on by García-Carpintero (2013), who argues that the two accounts are in tension. My account is more akin to the (file-less) mental *filing* account proposed by Goodman & Gray (2022), which posits that object-representations are individuated by epistemic relations on the input side *and* relations to behaviour on the output side. Such an account, and my proposed account, can also help overcome some of the mental file account's more general issues, for instance those concerning relational properties: if the tree is next to a rock, in which file do I store that information? According to my account, this problem doesn't arise: I instantiate two object-tokens in a cognitive map and the relation between the tokens gives their spatial arrangement.



my representation that represents someone to have a broken arm. I might then, for some reason or other, judge that I am identical with this person. I would now instantiate a self\*-token that attributes to me the property of having a broken arm. As this judgement of identity is fallible, the resulting self-ascription isn't IEM. We can compare this situation to one in which I use proprioceptive information to determine that my arm is broken. This judgement is IEM since it's based on the articulation of the sensorimotor representation's unarticulated subject component.

Note that at least *some* self\*-token values must derive from the articulation of an egocentric representation's unarticulated subject. A judgement of identity between a self\*-token and some object-token requires a self\*-token. Now, if this self\*-token's value is itself derived from a judgement of identity with another self\*-token, then that requires yet another self\*-token. This well-known infinite regress (Evans, 1982; Fuchs, 2017; MacKenzie, 2010; Peacocke, 2014; Shoemaker, 1968) is only stopped when we arrive at a self\*-token whose value derives from the unarticulated subject component of a sensorimotor representation.

To conclude, when a sensorimotor representation's subject component is used to determine the value of a self\*-token, the resulting self-ascription is IEM. If, in contrast, a self\*-token's value stems from an identification with some object-token, then the self-ascribed properties aren't IEM. The account's ability to explain these phenomena provides further evidence that coordinated non-egocentric representations are genuine self-representations.

## Substantive self-representation

What makes minimal self-representations relatively common is that they are, as the name implies, minimal; they don't demand that self-representation be about specific kinds of property. We might consequently worry that minimal self-representation says a little too little about what is important about how, say, adult neurotypical human beings self-represent. Here, I want to make some cursory remarks about how minimal self-representation relates to such, more substantive, notions of self-representation.

Musholt believes that self-representation requires 'acquisition and application of the first person concept' (Musholt, 2013, sec. 4). And an agent can only do so by representing one's own and others' *mental states* (Musholt, 2012). This is a substantive view of self-representation: it argues that representing a specific kind of property – here, mental properties – is necessary for self-representation.

Many other kinds of property have been advanced as necessary (or at least important in some privileged sense) for self-representation. Grush (2000), taking inspiration from Strawson (2011), writes that 'the subject/object distinction is the result of a cognizer's representation of space' (p. 62). Peacocke (2014) and Campbell (1999) think that spatial properties are a good start, but that temporal properties are also neces-

sary. The necessity of representing one's temporal properties is also one of the few things the literature on narrative selves can agree on (Goldie, 2012; Lamarque, 2004; Menary, 2008). Yet another group of writers focuses on representations of one's own body (Hohwy & Michael, 2017; Metzinger, 2003). Hohwy, for instance, aims to explain 'self-representation in terms of inferred hidden causes' (p. 374) where these 'deeply hidden causes stem from the organism itself' (p. 375).

If something like my account is correct, the disagreements between these various proposals can't be about what constitutes a genuine self-representation. After all, any coordinated non-egocentric representation, no matter its content, is a genuine self-representation. To me, this suggests that one of two things is the case. First, there may be a genuine disagreement about the kinds of property that a creature must represent for the content of the self-representation to be sufficiently rich to capture a self's essential properties. Alternatively, there may only be a surface disagreement since these accounts try to capture different kinds of – possibly complementary – self-representation.

Peacocke (2014) appears to think that a self essentially exemplifies temporal and spatial properties. A creature with merely a non-egocentric spatial representation only represents '*this place on the map is here*' (p. 30, emphasis in the original) and *de se* content requires representing one's trajectory in time. However if that *here* location links to sensorimotor information in the sense discussed for the self\*-token throughout this paper, then it *is* a self\*-token, and the resulting representation is a self-representation. What's left, then, seems to be the claim that selves are essentially temporally extended and that a genuine self-representation must capture this property.

Similarly, Musholt seems to think that a self essentially exemplifies mental properties, and that failure to represent these *qua* mental properties is a failure to represent a self *qua* self. She writes that 'for a subject to realize that other subjects are distinct beings with their own mental states, and as such are similar to herself [...] she [...] needs to understand that others possess first person information about their mental states which might differ from her own' (Musholt, 2012, p. 78). It is therefore only at around one-and-a-half years of age that a 'child begins to explicitly differentiate between self and other' (p. 79). Note how this argument appears to presuppose that the child must understand what kind of being she is, where this essentially involves having mental states. Without representing oneself as a bearer of mental states, the child fails to represent herself as she is and thus fails to self-represent.

If this is really what these accounts disagree about, then their disagreement is about selves' essential properties. In other words, the disagreements seem to draw on differing answers regarding the metaphysics of selfhood rather than questions particular to self-representation.

Alternatively, we may want to interpret the differences not as disagreements but as simply describing different kinds of self-representation. As seen throughout this

paper, creatures may minimally self-represent different kinds of property. A creature can, of course, also self-represent *a range* of properties of different kinds. Thus, while I used fruit bats to illustrate *spatial* self-representation, it seems likely that they also temporally self-represent – this would allow them to infer that certain motor commands will make it so that they arrive at a target location *at some moment in the future*.

The differences between the various accounts I introduced could then be seen as differences in the richness of the self-representations involved. Such differences are important, and depending on the use to which one wants to put one's concept of self-representation, we might prefer one over the other. Researchers in the empirical sciences should focus on exploring how different creatures have self-representations of different kinds of (ranges of) property, rather than investigating whether this or that creature exemplifies this or that property deemed essential by philosophers.

I want to make two quick notes about building more substantive self-representations from minimal ones. First, whatever properties may be represented by a substantive representation, they must be represented in coordinated non-egocentric representations. Otherwise, they fail to fulfil even the minimum criteria for genuine self-representation. For a fruit bat to self-represent temporal and spatial properties, it must realise a spatial and temporal non-egocentric representation.

However, and this is the second point, this isn't sufficient; the relevant non-egocentric representations also need to be integrated. The representations need to be coordinated with sensorimotor information such that sensorimotor information is inferred as a function of the content of the temporal *and* spatial representation. The fruit bat represents that *it is at some location now* because it uses information about temporal *and* spatial properties together to infer sensorimotor information. Without such integration, the bat might realise two self-representations, but there would be no (tacit) representation of the relevant self\*-tokens referring to the same subject.

Human beings are then distinguished from bats by having strongly integrated and high-dimensional self-representations. A self\*-token in such a representation self-ascribes a multitude of properties: in addition to spatial, temporal, bodily, and mental properties, we also represent social properties, our character traits, narrative links between events, and much much more. While the debates concerning our self-representations' level of integration is far from settled – see the disagreements between Lambie & Marcel (2002) and Dalglish & Power (2004) or between Ismael (2008) and Clark (2007) – my account makes it clear that the difference between our self-representations and those exemplified by bats and scrub-jays is one of degree.

## Conclusion

I've suggested that some organisms use coordinated non-egocentric representations to navigate space, keep track of temporal relations, control their body movements, and

monitor what other people know about the world. Such representations explicitly differentiate between properties attributed to the self and those attributed to others. Additionally, they contain a self\*-token that enables coordination with sensorimotor information, which directly links the representation to behaviour and sensation. The proposed account explains why some self-attributed properties are immune to error through misidentification whereas others are not.

Whenever a system represents a property with a coordinated non-egocentric representation, it realises a minimal self-representation. As we have seen, young infants and a diverse roster of non-human animals can realise representations of this type, making self-representation more common than many might have believed. By integrating minimal self-representations of various kinds of property, more substantive forms of self-representation – such as those displayed by neurotypical adult human beings – emerge.

## Acknowledgements

I want to thank Mark Sprevak for feedback on early versions of the ideas developed in this paper, Hadeel Naeem for reading many drafts as the work progressed, and members of the Philosophy of Cognitive Science reading group at LOGOS (Universitat de Barcelona) for detailed discussions.

## Author statements

**Financial Support / Funding statement:** The author was supported by the grant P500PH\_202829 / 1 of the Swiss National Research Foundation.

**Competing Interest statement:** The author hereby declares not to have any conflict of interests (financial or non-financial) regarding the present research.

## References

- Apperly, Ian A., & Butterfill, Stephen A. (2009). Do humans have two systems to track beliefs and belief-like states? *Psychological Review*, 116, 953–970. <http://dx.doi.org/10.1037/a0016923>
- Arcediano, Francisco, Escobar, Martha, & Miller, Ralph R. (2003). Temporal integration and temporal backward associations in human and nonhuman subjects. *Animal Learning & Behavior*, 31, 242–256. <http://dx.doi.org/10.3758/bf03195986>
- Babb, Matthew. (2016). The essential indexicality of intentional action. *The Philosophical Quarterly*, 66, 439–457. <http://dx.doi.org/10.1093/pq/pqw023>
- Baciadonna, Luigi, Jerwood, Georgia M., Farrar, Benjamin G., Clayton, Nicola S., & Emery, Nathan J. (2022). Investigation of mirror-self recognition in ravens (*corvus*)

- corax). *Journal of Comparative Psychology*, 136, 194–198. <http://dx.doi.org/10.1037/com0000319>
- Baranes, Adrien, & Oudeyer, Pierre-Yves. (2013). Active learning of inverse models with intrinsically motivated goal exploration in robots. *Robotics and Autonomous Systems*, 61, 49–73. <http://dx.doi.org/10.1016/j.robot.2012.05.008>
- Barnet, Robert C., Cole, Robert P., & Miller, Ralph R. (1997). Temporal integration in second-order conditioning and sensory preconditioning. *Animal Learning & Behavior*, 25, 221–233. <http://dx.doi.org/10.3758/bf03199061>
- Bermúdez, Jose Luis. (1998). *The paradox of self-consciousness*. MIT Press.
- Bermúdez, José Luis. (2017). Yes, essential indexicals really are essential\*. *Analysis*, 77, 690–694. <http://dx.doi.org/10.1093/analys/anx137>
- Berthier, Neil E. (1996). Learning to reach: A mathematical model. *Developmental Psychology*, 32, 811–823. <http://dx.doi.org/10.1037/0012-1649.32.5.811>
- Berthier, Neil E., & Keen, Rachel. (2006). Development of reaching in infancy. *Experimental Brain Research*, 169, 507–518. <http://dx.doi.org/10.1007/s00221-005-0169-9>
- Bruner, J. S., & Koslowski, Barbara. (1972). Visually preadapted constituents of manipulatory action. *Perception*, 1, 3–14. <http://dx.doi.org/10.1068/p010003>
- Bugnyar, Thomas, Reber, Stephan A., & Buckner, Cameron. (2016). Ravens attribute visual access to unseen competitors. *Nature Communications*, 7. <http://dx.doi.org/10.1038/ncomms10506>
- Buhusi, Catalin V., & Meck, Warren H. (2005). What makes us tick? Functional and neural mechanisms of interval timing. *Nature Reviews Neuroscience*, 6, 755–765. <http://dx.doi.org/10.1038/nrn1764>
- Burge, Tyler. (2018). Do infants and nonhuman animals attribute mental states? *Psychological Review*, 125, 409–434. <http://dx.doi.org/10.1037/rev0000091>
- Butterfill, Stephen A., & Apperly, Ian A. (2013). How to construct a minimal theory of mind. *Mind & Language*, 28, 606–637. <http://dx.doi.org/10.1111/mila.12036>
- Campbell, John. (1999). *Past, space, and self*. NetLibrary, Inc.
- Cappelen, Herman, & Dever, Josh. (2014). *The inessential indexical: On the philosophical insignificance of perspective and the first person* (1 edition). Oxford University Press.
- Castañeda, Hector-Neri. (1966). 'He': A study in the logic of self-consciousness. *Ratio*, 8(December), 130–157.
- Castellini, Claudio, Orabona, Francesco, Metta, Giorgio, & Sandini, Giulio. (2007). Internal models of reaching and grasping. *Advanced Robotics*, 21, 1545–1564. <http://dx.doi.org/10.1163/156855307782148587>
- Clark, Andy. (2007). Soft selves and ecological control. In *Distributed cognition and the will: Individual volition and social context* (pp. 101–122). MIT Press.
- Clayton, Nicola S., Dally, Joanna M., & Emery, Nathan J. (2007). Social cognition by food-caching corvids. The western scrub-jay as a natural psychologist. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 362, 507–522. <http://dx.doi.org/10.1098/rstb.2006.2200>

- [//dx.doi.org/10.1098/rstb.2006.1992](http://dx.doi.org/10.1098/rstb.2006.1992)
- Clifton, Rachel K., Muir, Darwin W., Ashmead, Daniel H., & Clarkson, Marsha G. (1993). Is visually guided reaching in early infancy a myth? *Child Development*, 64, 1099–1110. <http://dx.doi.org/10.1111/j.1467-8624.1993.tb04189.x>
- Collett, Thomas S. (2019). Path integration: How details of the honeybee waggle dance and the foraging strategies of desert ants might help in understanding its mechanisms. *Journal of Experimental Biology*, 222. <http://dx.doi.org/10.1242/jeb.205187>
- Corbetta, Daniela, Thurman, Sabrina L., Wiener, Rebecca F., Guan, Yu, & Williams, Joshua L. (2014). Mapping the feel of the arm with the sight of the object: On the embodied origins of infant reaching. *Frontiers in Psychology*, 5. <http://dx.doi.org/10.3389/fpsyg.2014.00576>
- Corbetta, Daniela, Wiener, Rebecca F., Thurman, Sabrina L., & McMahon, Emalie. (2018). The embodied origins of infant reaching: Implications for the emergence of eye-hand coordination. *Kinesiology Review*, 7, 10–17. <http://dx.doi.org/10.1123/k.r.2017-0052>
- Coslett, H. Branch, Buxbaum, Laurel J., & Schwobel, John. (2008). Accurate reaching after active but not passive movements of the head: Evidence for forward modeling. *Behavioural Neurology*, 19, 117–125.
- Cothi, William de, Nyberg, Nils, Griesbauer, Eva-Maria, Ghanamé, Carole, Zisch, Fiona, Lefort, Julie M., Fletcher, Lydia, Newton, Coco, Renaudineau, Sophie, Bendor, Daniel, Grieves, Roddy, Duvelle, Éléonore, Barry, Caswell, & Spiers, Hugo J. (2022). Predictive maps in rats and humans for spatial navigation. *Current Biology*, 32, 3676–3689.e5. <http://dx.doi.org/10.1016/j.cub.2022.06.090>
- Dalglish, Tim, & Power, Michael J. (2004). The i of the storm—relations between self and conscious emotion experience: Comment on lambie and marcel (2002). *Psychological Review*, 111, 812–818. <http://dx.doi.org/10.1037/0033-295x.111.3.812>
- Dennett, Daniel C. (1982). Styles of mental representation. *Proceedings of the Aristotelian Society*, 83, 213–226. <https://www.jstor.org/stable/4545000>
- Desmurget, Michel, & Grafton, Scott. (2000). Forward modeling allows feedback control for fast reaching movements. *Trends in Cognitive Sciences*, 4, 423–431. [http://dx.doi.org/10.1016/s1364-6613\(00\)01537-0](http://dx.doi.org/10.1016/s1364-6613(00)01537-0)
- Dienes, Zoltan, & Perner, Josef. (1999). A theory of implicit and explicit knowledge. *Behavioral and Brain Sciences*, 22, 735–808. <http://dx.doi.org/10.1017/s0140525x99002186>
- Emery, N. J., & Clayton, N. S. (2001). Effects of experience and social context on prospective caching strategies by scrub jays. *Nature*, 414, 443–446. <http://dx.doi.org/10.1038/35106560>
- Emery, Nathan J., & Clayton, Nicola S. (2008). How to build a scrub-jay that reads minds. In Shoji Itakura & Kazuo Fujita (Eds.), *Origins of the social mind: Evolutionary and developmental views*. Springer Japan.
- Evans, Gareth. (1982). *The varieties of reference* (John Henry McDowell, Ed.). Claren-

- don Press; Oxford University Press.
- Filimon, Flavia. (2015). Are all spatial reference frames egocentric? Reinterpreting evidence for allocentric, object-centered, or world-centered reference frames. *Frontiers in Human Neuroscience*, 9. <http://dx.doi.org/10.3389/fnhum.2015.00648>
- Fuchs, Thomas. (2017). Self across time: The diachronic unity of bodily existence. *Phenomenology and the Cognitive Sciences*, 16(2), 291–315.
- Gallistel, C. R. (1989). Animal cognition: The representation of space, time and number. *Annual Review of Psychology*, 40(1), 155–189. <https://doi.org/10.1146/annurev.ps.40.020189.001103>
- García-Carpintero, Manuel. (2013). The self file and immunity to error through misidentification. *Disputatio*, 5, 191–206. <http://dx.doi.org/10.2478/disp-2013-0017>
- Goldie, Peter. (2012). *The mess inside: Narrative, emotion, and the mind*. Oxford University Press.
- Goodman, Rachel, & Gray, Aidan. (2022). Mental filing. *Noûs*, 56, 204–226. <http://dx.doi.org/10.1111/nous.12354>
- Grush, Rick. (2000). Self, world and space: The meaning and mechanisms of ego- and allocentric spatial representation. *Brain and Mind*, 1(1), 59–92.
- Grush, Rick. (2004). The emulation theory of representation: Motor control, imagery, and perception. *Behavioral and Brain Sciences*, 27, 377–396. <http://dx.doi.org/10.1017/S0140525X04000093>
- Grush, Rick. (2008). Temporal representation and dynamics. *New Ideas in Psychology*, 26(2), 146–157. <http://www.sciencedirect.com/science/article/pii/S0732118X07000475>
- Gür, Ezgi, Duyan, Yalçın Akın, & Balcı, Fuat. (2018). Spontaneous integration of temporal information: Implications for representational/computational capacity of animals. *Animal Cognition*, 21, 3–19. <http://dx.doi.org/10.1007/s10071-017-1137-z>
- Hofsten, Claes von. (1982). Eye–hand coordination in the newborn. *Developmental Psychology*, 18, 450–461. <http://dx.doi.org/10.1037/0012-1649.18.3.450>
- Hofsten, Claes von. (1991). Structuring of early reaching movements: A longitudinal study. *Journal of Motor Behavior*, 23, 280–292. <http://dx.doi.org/10.1080/00222895.1991.9942039>
- Hofsten, Claes von, & Fazel-Zandy, Shirin. (1984). Development of visually guided hand orientation in reaching. *Journal of Experimental Child Psychology*, 38, 208–219. [http://dx.doi.org/10.1016/0022-0965\(84\)90122-x](http://dx.doi.org/10.1016/0022-0965(84)90122-x)
- Hofsten, Claes von, & Lindhagen, Karin. (1979). Observations on the development of reaching for moving objects. *Journal of Experimental Child Psychology*, 28, 158–173. [http://dx.doi.org/10.1016/0022-0965\(79\)90109-7](http://dx.doi.org/10.1016/0022-0965(79)90109-7)
- Hohwy, Jakob, & Michael, John. (2017). Why should any body have a self? In Frédérique de Vignemont & Adrian J. T. Alsmith (Eds.), *The subject's matter: Self-consciousness and the body* (pp. 363–391). The MIT Press.
- Ismael, Jenann. (2008). *Situated self*. Oxford University Press, Incorporated.

- Ismael, Jenann. (2012). Immunity to error as an artefact of transition between representational media. In *Immunity to error through misidentification: New essays* (pp. 62–80). Cambridge University Press.
- Klatzky, Roberta L. (1998). Allocentric and egocentric spatial representations: Definitions, distinctions, and interconnections. In *Lecture Notes in Computer Science* (pp. 1–17). Springer Berlin Heidelberg. [http://dx.doi.org/10.1007/3-540-69342-4\\_1](http://dx.doi.org/10.1007/3-540-69342-4_1)
- Konczak, Jürgen, Borutta, Maïke, Topka, Helge, & Dichgans, Johannes. (1995). The development of goal-directed reaching in infants: Hand trajectory formation and joint torque control. *Experimental Brain Research*, 106. <http://dx.doi.org/10.1007/bf00241365>
- Körding, Konrad P., & Wolpert, Daniel M. (2006). Bayesian decision theory in sensorimotor control. *Trends in Cognitive Sciences*, 10(7), 319–326. <http://linkinghub.elsevier.com/retrieve/pii/S1364661306001276>
- Kort, Selvino R. de, Dickinson, Anthony, & Clayton, Nicola S. (2005). Retrospective cognition by food-caching western scrub-jays. *Learning and Motivation*, 36, 159–176. <http://dx.doi.org/10.1016/j.lmot.2005.02.008>
- Lamarque, Peter. (2004). On not expecting too much from narrative. *Mind and Language*, 19(4), 393–408. <http://doi.wiley.com/10.1111/j.0268-1064.2004.00265.x>
- Lambie, John A., & Marcel, Anthony J. (2002). Consciousness and the varieties of emotion experience: A theoretical framework. *Psychological Review*, 109, 219–259. <http://dx.doi.org/10.1037//0033-295x.109.2.219>
- Lepage, Jean-François, & Théoret, Hugo. (2007). The mirror neuron system: Grasping others? Actions from birth? *Developmental Science*, 10, 513–523. <http://dx.doi.org/10.1111/j.1467-7687.2007.00631.x>
- Lewis, David. (1979). Attitudes de dicto and de se. *The Philosophical Review*, 88, 513. <http://dx.doi.org/10.2307/2184843>
- Mach, Ernst. (1890). The analysis of the sensations. antimetaphysical. *The Monist*, 1(1), 48–68. <http://www.jstor.org/stable/27896829>
- MacKenzie, Matthew. (2010). Enacting the self: Buddhist and enactivist approaches to the emergence of the self. *Phenomenology and the Cognitive Sciences*, 9(1), 75–99.
- Marshall, Peter J., & Meltzoff, Andrew N. (2015). Body maps in the infant brain. *Trends in Cognitive Sciences*, 19, 499–505. <http://dx.doi.org/10.1016/j.tics.2015.06.012>
- McCormack, Teresa. (2015). The development of temporal cognition. *Handbook of Child Psychology and Developmental Science*, 1–47. <http://dx.doi.org/10.1002/9781118963418.childpsy215>
- Meltzoff, Andrew N., Saby, Joni N., & Marshall, Peter J. (2019). Neural representations of the body in 60-day-old human infants. *Developmental Science*, 22. <http://dx.doi.org/10.1111/desc.12698>
- Menary, Richard. (2008). Embodied narratives. *Journal of Consciousness Studies*, 15(6), 63–84.



- Metzinger, Thomas. (2003). *Being no one: The self-model theory of subjectivity*. MIT Press.
- Millikan, Ruth Garrett. (1990). The myth of the essential indexical. *Noûs*, 24, 723. <http://dx.doi.org/10.2307/2215811>
- Molet, Mikael, Miguez, Gonzalo, Cham, Henry X., & Miller, Ralph R. (2012). When does integration of independently acquired temporal relationships take place? *Journal of Experimental Psychology: Animal Behavior Processes*, 38, 369–380. <http://dx.doi.org/10.1037/a0029379>
- Musholt, Kristina. (2012). Self-consciousness and intersubjectivity. *Grazer Philosophische Studien*, 84(1), 63–89.
- Musholt, Kristina. (2013). Self-consciousness and nonconceptual content. *Philosophical Studies*, 163(3), 649–672. <http://link.springer.com/10.1007/s11098-011-9837-8>
- Musholt, Kristina. (2015). *Thinking about oneself: From nonconceptual content to the concept of a self*. MIT Press.
- Nguyen, Phuong D. H., Eppe, Manfred, & Wermter, Stefan. (2020). *Robotic self-representation improves manipulation skills and transfer learning*. <http://arxiv.org/abs/2011.06985v1>
- O'Keefe, John, & Burgess, Neil. (1996). Geometric determinants of the place fields of hippocampal neurons. *Nature*, 381, 425–428. <http://dx.doi.org/10.1038/381425a0>
- O'Keefe, John, & Nadel, Lynn. (1978). *The hippocampus as a cognitive map*. Clarendon Press ; Oxford University Press.
- Peacocke, Christopher. (2014). *The mirror of the world: Subjects, consciousness, and self-consciousness*. Oxford University Press.
- Penn, Derek C., Holyoak, Keith J., & Povinelli, Daniel J. (2008). Darwin's mistake: Explaining the discontinuity between human and nonhuman minds. *Behavioral and Brain Sciences*, 31, 109–130. <http://dx.doi.org/10.1017/S0140525X08003543>
- Penn, Derek C., & Povinelli, Daniel J. (2007). On the lack of evidence that non-human animals possess anything remotely resembling a 'theory of mind'. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 362, 731–744. <http://dx.doi.org/10.1098/rstb.2006.2023>
- Perry, John. (1979). The problem of the essential indexical. *Noûs*, 13(1), 3–21.
- Perry, John. (1993). *The problem of the essential indexical: And other essays*. Oxford University Press.
- Perry, John. (2010). Selves and self-concepts. In Joseph Keim Campbell, Michael O'Rourke, & Harry Silverstein (Eds.), *Time and identity* (pp. 229–247). MIT Press.
- Perry, John, & Blackburn, Simon. (1986). Thought without representation. *Aristotelian Society Supplementary Volume*, 60, 137–166. <http://dx.doi.org/10.1093/aristoteliansupp/60.1.137>
- Prosser, Simon. (2015). Why are indexicals essential? *Proceedings of the Aristotelian Society*, 115, 211–233. <http://dx.doi.org/10.1111/j.1467-9264.2015.00392.x>
- Recanati, François. (2009). De re and de se. *Dialectica*, 63, 249–269. <http://dx.doi.org>

- /10.1111/j.1746-8361.2009.01194.x
- Recanati, François. (2012a). Immunity to error through misidentification: What it is and where it comes from. In *Immunity to error through misidentification: New essays* (pp. 180–201). Cambridge University Press.
- Recanati, François. (2012b). *Mental files*. Oxford University Press.
- Recanati, François. (2024). IEM explained. *Philosophical Psychology*, 1–20. <http://dx.doi.org/10.1080/09515089.2024.2359493>
- Reid, Chris R., Latty, Tanya, Dussutour, Audrey, & Beekman, Madeleine. (2012). Slime mold uses an externalized spatial ‘memory’ to navigate in complex environments. *Proceedings of the National Academy of Sciences*, 109(43), 17490–17494. <https://www.pnas.org/content/109/43/17490>
- Rescorla, Michael. (2017). Maps in the head? In *The routledge handbook of philosophy of animal minds*. Routledge.
- Rolf, Matthias, & Steill, Jochen J. (2012). Goal babbling: A new concept for early sensorimotor exploration. *Proceedings of Workshop on Developmental Robotics*.
- Saigusa, Tetsu, Tero, Atsushi, Nakagaki, Toshiyuki, & Kuramoto, Yoshiki. (2008). Amoebae anticipate periodic events. *Physical Review Letters*, 100(1), 018101. <https://link.aps.org/doi/10.1103/PhysRevLett.100.018101>
- Schillaci, Guido, Hafner, Verena V., & Lara, Bruno. (2016). Exploration behaviors, body representations, and simulation processes for the development of cognition in artificial agents. *Frontiers in Robotics and AI*, 3. <http://dx.doi.org/10.3389/frobt.2016.00039>
- Searle, John R. (1983). *Intentionality: An essay in the philosophy of mind*. Cambridge University Press.
- Shoemaker, Sydney. (1968). Self-reference and self-awareness. *Journal of Philosophy*, 65(October), 555–567.
- Slaughter, Virginia, & Brownell, Celia A. (Eds.). (2011). *Early development of body representations*.
- Srinivasan, Mandyam V. (2015). Where paths meet and cross: Navigation by path integration in the desert ant and the honeybee. *Journal of Comparative Physiology A*, 201, 533–546. <http://dx.doi.org/10.1007/s00359-015-1000-0>
- Stachenfeld, Kimberly L., Botvinick, Matthew M., & Gershman, Samuel J. (2017). The hippocampus as a predictive map. *Nature Neuroscience*, 20, 1643–1653. <http://dx.doi.org/10.1038/nn.4650>
- Stalnaker, Robert C. (1981). Indexical belief. *Synthese*, 49, 192–151.
- Stoewer, Paul, Schlieker, Christian, Schilling, Achim, Metzner, Claus, Maier, Andreas, & Krauss, Patrick. (2022). Neural network based successor representations to form cognitive maps of space and language. *Scientific Reports*, 12. <http://dx.doi.org/10.1038/s41598-022-14916-1>
- Stoytchev, A. (2009). Some basic principles of developmental robotics. *IEEE Transactions on Autonomous Mental Development*, 1, 122–130. <http://dx.doi.org/10.1109/t>

amd.2009.2029989

- Strawson, Peter Frederick. (2011). *Individuals: An essay in descriptive metaphysics* (Reprinted, transferred to digital printing). Routledge.
- Takemura, Naohiro, Inui, Toshio, & Fukui, Takao. (2018). A neural network model for development of reaching and pointing based on the interaction of forward and inverse transformations. *Developmental Science*, 21, e12565. <http://dx.doi.org/10.1111/desc.12565>
- Taylor, Kathleen M., Joseph, Victory, Zhao, Alice S., & Balsam, Peter D. (2014). Temporal maps in appetitive pavlovian conditioning. *Behavioural Processes*, 101, 15–22. <http://dx.doi.org/10.1016/j.beproc.2013.08.015>
- Thelen, Esther, Corbetta, Daniela, & Spencer, John P. (1996). Development of reaching during the first year: Role of movement speed. *Journal of Experimental Psychology: Human Perception and Performance*, 22, 1059–1076. <http://dx.doi.org/10.1037/0096-1523.22.5.1059>
- Tolman, Edward C. (1948). Cognitive maps in rats and men. *Psychological Review*, 55, 189–208. <http://dx.doi.org/10.1037/h0061626>
- Tsoar, Asaf, Nathan, Ran, Bartan, Yoav, Vyssotski, Alexei, Dell’Omo, Giacomo, & Ulanovsky, Nachum. (2011). Large-scale navigational map in a mammal. *Proceedings of the National Academy of Sciences*, 108. <http://dx.doi.org/10.1073/pnas.1107365108>
- Wehner, Rüdiger, & Srinivasan, Mandyam V. (1981). Searching behaviour of desert ants, genus *cataglyphis* (formicidae, hymenoptera). *Journal of Comparative Physiology A*, 142, 315–338. <http://dx.doi.org/10.1007/bf00605445>
- Wentworth, Naomi, Benson, Janette B., & Haith, Marshall M. (2000). The development of infants’ reaches for stationary and moving targets. *Child Development*, 71, 576–601. <http://dx.doi.org/10.1111/1467-8624.00169>
- Witherington, David C. (2005). The development of prospective grasping control between 5 and 7 months: A longitudinal study. *Infancy*, 7, 143–161. [http://dx.doi.org/10.1207/s15327078in0702\\_2](http://dx.doi.org/10.1207/s15327078in0702_2)
- Wittgenstein, Ludwig. (2007). *The blue and brown books: Preliminary studies for the "philosophical investigations"*. Blackwell.
- Wittlinger, Matthias, Wehner, Rüdiger, & Wolf, Harald. (2006). The ant odometer: Stepping on stilts and stumps. *Science*, 312, 1965–1967. <http://dx.doi.org/10.1126/science.1126912>
- Wolpert, Daniel M., & Ghahramani, Zoubin. (2000). Computational principles of movement neuroscience. *Nature Neuroscience*, 3, 1212–1217. <http://dx.doi.org/10.1038/81497>