

# Self-consciousness and nonconceptual content

Kristina Musholt

© Springer Science+Business Media B.V. 2011

**Abstract** Self-consciousness can be defined as the ability to think 'I'-thoughts. Recently, it has been suggested that self-consciousness in this sense can (and should) be accounted for in terms of nonconceptual forms of self-representation. Here, I will argue that while theories of nonconceptual self-consciousness do provide us with important insights regarding the essential genetic and epistemic features of self-conscious thought, they can only deliver part of the full story that is required to understand the phenomenon of self-consciousness. I will provide two arguments to this effect, drawing on insights from the philosophy of language and on structural differences between conceptual and nonconceptual forms of representation. Both arguments rest on the intuition that while self-consciousness requires explicit self-representation, nonconceptual content can at best provide implicit self-related information. I will conclude that in order to explain the emergence of self-conscious thought out of more basic forms of representation one has to explain the transition between implicit self-related information and explicit self-representation.

**Keywords** Self-consciousness · Consciousness · Nonconceptual content · Implicit · Explicit · Representation · Immunity to error through misidentification · Unarticulated constituents · Self-reference · Perception · Proprioception

## 1 Self-consciousness

Self-consciousness can be defined as the ability to think 'I'-thoughts.<sup>1</sup> 'I'-thoughts are thoughts with first-person content that non-accidentally refer to oneself. In other

---

<sup>1</sup> Cf., for example, Bermúdez (1998, 2001) and Vosgerau (2009).

---

K. Musholt (✉)  
Department of Philosophy, Logic and Scientific Method, London School of Economics and Political  
Science, Houghton Street, London WC2A 2AE, UK  
e-mail: K.Musholt@lse.ac.uk

words, when entertaining these thoughts, the subject is (necessarily) aware of the self-referring nature of the thought; it is not possible to entertain an ‘I’-thought without knowing that it refers to oneself. Consider this well-known story by John Perry:

I once followed a trail of sugar on a supermarket floor, pushing my cart down the aisle on one side of a tall counter and back the aisle on the other, seeking the shopper with the torn sack to tell him he was making a mess. With each trip around the counter, the trail became thicker. But I seemed unable to catch up. Finally it dawned on me. I was the shopper I was trying to catch (Perry 2000, p. 27).

What is happening here? Obviously, Perry was aware from the start that *someone* was making a mess, and this thought did de facto refer to him. Nevertheless, he was not aware that the thought referred to him until he realized that he himself was making the mess. What Perry is after here is the fact that one can entertain many thoughts that are in fact about oneself, such as, in his example, ‘Someone is making a mess’, ‘The only bearded philosopher in the market is making a mess’, or even ‘John Perry is making a mess’ without realizing that these thoughts refer to oneself (one might suffer from amnesia and hence might not remember one’s name, for example). In contrast to these *de re* beliefs about oneself, there is a class of thoughts, namely *de se* thoughts, that can only be entertained by a subject if that subject is aware that the thoughts are self-referring. In other words, these thoughts are non-accidentally self-referring. The canonical expression of these thoughts involves the first-person pronoun. Perry can only come to think ‘I am making a mess!’ if he realizes that it is in fact he himself who is carrying a leaking sugar bag, for in order to competently employ the first-person pronoun he has to know that it necessarily refers to him when he is entertaining this thought. And it is this, non-accidentally self-referring, thought that will prompt him to stop the cart and fix the leak.

So in contrast to *de re* or *de dicto*<sup>2</sup> thoughts about myself, *de se* thoughts are thoughts that are known by me to be about myself, and that have immediate relevance for action. These thoughts are expressed in direct speech by means of the first-person pronoun (and, following Castañeda, in indirect speech by means of *s/he\**).

However, as Wittgenstein (1958) pointed out, upon closer examination we can actually distinguish between two different uses of the first-person pronoun: the use-as-subject and the use-as-object. Examples for the latter are thoughts such as ‘My arm is broken’, ‘I have grown six inches’, or ‘I have a stain on my shirt’. Examples for the first, where ‘I’ is used as subject, are ‘I see a tree’ or ‘I hear the birds outside my window’. If I use the first-person pronoun ‘as-object’, as in ‘I have a stain on my shirt’ it is theoretically possible that I am mistaken with regard to the question of *who* has a stained shirt. For example, I might be seeing a stained shirt in the mirror, mistaking it for my own, while in fact it is my neighbor’s, very similar looking,

<sup>1</sup> Cf., for example, Bermúdez (1998, 2001) and Vosgerau (2009).

<sup>2</sup> Although *de re* and *de dicto* thoughts are not to be confused with each other, we can ignore the distinction between them for the purposes of this article. For a detailed discussion see (Castañeda et al. 1999; Chisholm 1981; Frank 2007).

shirt. In this situation, I am committing an error of misidentification when claiming on the basis of seeing someone else's shirt that I have a stained shirt. Contrary to this scenario, I can principally not be mistaken with regard to the identity of the subject when I say (on the basis of my visual experience) 'I see a tree'. In this case, I could be mistaken about the fact that it is a tree that I see (perhaps I am mistaking it for a different object), but I cannot possibly be mistaken about the fact that it is *me* who has this (possibly non-veridical) perception. In other words, in this latter case, I cannot commit an error of misidentification relative to the first-person pronoun. Shoemaker introduces the notion of immunity to error through misidentification in the following, well-known way:

To say that a statement "a is  $\phi$ " is subject to error through misidentification relative to the term 'a' means that the following is possible: the speaker knows some particular thing to be  $\phi$ , but makes the mistake of asserting "a is  $\phi$ " because, and only because, he mistakenly thinks that the thing he knows to be  $\phi$  is what 'a' refers to (Shoemaker 1968, p. 557).

The kind of mistake described in this quote is not possible in statements that are immune to error through misidentification with respect to the term 'a'. So while a self-conscious thought that is immune to error through misidentification can misrepresent the property that is being ascribed, it cannot misrepresent the subject purportedly possessing that property. The reason for this, according to Wittgenstein and Shoemaker, is that these self-ascriptions do not involve the identification of a subject. If no identification of a subject is involved, it is impossible to misidentify it. This is why Wittgenstein terms the use of the first-person pronoun involved in these kinds of first-person judgments use-as-subject as opposed to use-as-object. When making a first-person judgment that is immune to error through misidentification we are expressing a type of awareness of ourselves not as objects (that have to be identified), but as subjects. It is this kind of awareness that is essential to the phenomenon of self-consciousness, as it seems to indicate a particular (immediate) way of gaining knowledge about ourselves that is different in principle from the ways in which we can gain knowledge about others.<sup>3</sup>

Hence, 'I'-thoughts (where the first person pronoun is used 'as-subject') can be characterized by three essential features: (1) non-accidental self-reference, (2) immediate action-relevance, and (3) immunity to error through misidentification. A theory of self-consciousness must account for all of these features.

## 2 Self-consciousness and nonconceptual content

As we have seen, linguistically, *de se* thoughts are canonically expressed by means of the first person pronoun, such as in 'I  $\phi$ ' where  $\phi$  stands for some predicate (e.g.,

<sup>3</sup> It is important to note that what makes a judgment immune to error through misidentification is the evidence on which it is based, not the subject matter or the fact that it employs the first-person pronoun (Evans 1982, Ch. 7).

<sup>4</sup> Or, in indirect speech, by means of the indirect reflexive pronoun, commonly marked as 'he\*' (as in 'he himself') (Castaneda 1966).

in the utterance ‘I am thirsty’, ‘I feel pain’, etc.).<sup>4</sup> Consequently, many traditional philosophical theories of self-consciousness assume, following Dummett’s view that the analysis of thought must proceed via the analysis of language, that since (a) self-consciousness can be characterized by ‘I’-thoughts, and (b) I-thoughts can be characterized by their canonical linguistic expression, all we need to do in order to understand self-consciousness is to analyze the semantics of the first person pronoun (Dummett 1973).

However, more recently, attempts have been made to give an account of the essential features of full-blown self-consciousness in terms of simpler, more primitive, and notably non-linguistic and nonconceptual forms of self-representation (Bermúdez 1998; Hurley 1997; Vosgerau 2009).<sup>5</sup> In other words, it has been suggested that the ability to entertain ‘I’- or *de se* thoughts is independent of the mastery of the first person pronoun and of concept possession and, moreover, that we need to appeal to certain types of nonconceptual representation in order to explain the genetic origins as well as the epistemic features of self-conscious thoughts.<sup>6</sup>

The attempt to show how full-blown self-consciousness emerges out of simpler forms of representations is commendable for several reasons. First, the analysis of self-conscious thought in terms of an analysis of the semantics of the first person pronoun seems to lead into what has been called ‘the paradox of self-consciousness’ (Bermúdez 1998; Bermúdez 2001). According to Bermúdez, this paradox is constituted by two forms of vicious circularity, an explanatory circularity and a capacity circularity. The first type of circularity arises because one cannot give an account of the mastery of the first person pronoun without referring to the capacity to think ‘I’-thoughts. The second type of circularity arises because, given the interdependence of the capacity for self-conscious thought and mastery of the first person pronoun, it is impossible to explain how self-consciousness develops in the first place (cf. Bermúdez 2001, pp. 131–32). Whether or not one agrees that this problem reaches the status of a paradox, it seems uncontroversial that self-consciousness is an ability that develops over the course of phylogeny and ontogeny, and that any plausible philosophical theory of self-consciousness should therefore at the very least allow for an explanation of, and at best be able to give an account of this development. Second, such an account would help us answer the question of how much of our mental life is shared with non-linguistic animals.

However, the notion of nonconceptual self-consciousness is problematic, or so I will argue. One of the most prominent and best developed proposals of how to explain the ability for self-conscious thought independent of conceptual abilities can be found in Bermúdez (1998, 2001). Therefore, my discussion will mostly focus on

<sup>5</sup> For phenomenologically oriented attempts to give an account of self-consciousness in terms of basic, nonconceptual forms of self-representation also see, for example, Frank (2007); Legrand (2007); Poellner (2003); Zahavi (2005). These authors argue that full-fledged, reflective self-consciousness must ultimately be based on a pre-reflective, immediate familiarity with oneself, a sense of *mineness*, which accompanies all conscious experience. However, it is far from obvious how explanatory the notion of a pre-reflective self-consciousness really is, how it is to be spelled out in positive terms, and whether it truly reflects an irreducible aspect of our conscious experience (Schear 2009). In the interest of not further complicating my argumentation, I will not discuss this notion here.

<sup>6</sup> See Bermúdez (1998, 2002) for an explicit discussion of the relation between genetic and epistemic sources of self-consciousness.

his theory (although my arguments will apply to other accounts of nonconceptual self-consciousness as well).

## 2.1 Nonconceptual content

Nonconceptual content is standardly defined as content that can be ascribed to a creature despite that creature lacking the concepts required to specify that content. There are many good reasons to assume the existence of nonconceptual content in general. For instance, many non- and pre-linguistic beings, such as animals and young infants, display flexible behavior<sup>7</sup> the explanation of which seems to require the notion of representational content, despite the fact that we would not want to ascribe the possession of concepts to these beings (Bermúdez 1998; Hurley 2006; Peacocke 2002). Moreover, concept acquisition cannot be explained without assuming that there are nonconceptual forms of representation (Roskies 2008). Other arguments in favor of nonconceptual content refer to the fact that our perceptual experience seems far more detailed and that the perceptual discriminations we are able to make are much more fine-grained than the concepts we possess (see, for example, Heck 2000).

Nonconceptual content is standardly defined negatively, in opposition to conceptual content. While there is no general agreement regarding the correct theory of concepts, it is generally agreed among members of the debate between conceptualism and nonconceptualism that conceptual content, at the very least, consists of several components that can be decomposed and recombined. In other words, conceptual content is generally considered to meet Evans' Generality Constraint.<sup>8</sup> Nonconceptual content, then, is not subject to the Generality Constraint. We might also say that it is structure-implicit, which is to say that it contains no explicit representation of subject/objects and predicates (Dienes and Perner 1999).

This leaves open how nonconceptual content should be characterized in positive terms. While an answer to this question is clearly beyond the scope of this article, let me briefly sketch one view, which I am sympathetic to. On this view, nonconceptual content presents the world not in terms of truth-conditions (as conceptual content would do), but in terms of the possible interactions of the organism with its environment (Cussins 1990; Pettit 2003; Ward et al. 2010). So on this view, nonconceptual content presents the subject not with identifiable objects and their properties as such, but rather with possibilities for intentional action, where these include so-called 'epistemic actions', such as grouping, sorting and tracking objects. For instance, Pettit (2003) suggests that an object looks red insofar as "it manifestly enables you to sift and sort and track it in the red-appropriate manner" and the "ball that someone throws looks to be going fast so far as it manifestly elicits reaching *there* if I am to catch it, or ducking *now* if I am to avoid it" (p. 230; italics original). Similarly, according to Ward et al. (2010), it is the implicit knowledge of how to

<sup>7</sup> That is, behaviour that shows no law-like relation between sensory input and behavioural output.

<sup>8</sup> The basic idea is that conceptual mastery of the thought 'a is F' implies the ability to think 'a is G' for any property G of which the subject of the thought has a conception, and similarly to think 'b is F' for any object b of which she has a conception (Evans 1982, p. Sec. 4.3).

pursue and accomplish one's goals and intentions with regard to certain objects in one's environment that determines the content of one's perceptual experience of this environment. For instance, my experience of the mug in front of me can be captured by referring to my abilities to grasp the mug, to track it through space and time, and to my sensitivities towards changes in the mug's features (cf. Cussins 1990). Another way of putting this is by saying that nonconceptual content presents the world to the subject in terms of affordances, to employ a Gibsonian notion.<sup>9</sup> We might also express this by saying that nonconceptual representations should be understood in terms of 'knowledge-how' rather than 'knowledge-that'.<sup>10</sup> When these nonconceptual representations are conceptualized they will underwrite the subject's judgment that there is an object with certain properties (which can be identified as a green tree, for example), but the nonconceptual perceptual representations are not themselves structured in this way.

As we will see in the following, this view sits well with the idea that perceptual content necessarily contains self-related information, but it does not sit well with the idea that nonconceptual content ought to be characterized as (proto-) compositional content.<sup>11</sup>

## 2.2 Nonconceptual self-consciousness

We have just seen that nonconceptual content is content that can be ascribed to a creature despite that creature lacking the concepts required to specify that content. Nonconceptual self-consciousness, then, is *first-person* content that can be ascribed to a creature despite the fact that this creature lacks the concepts required to specify this content (in particular a self-concept). And in order to qualify as a form of genuine self-consciousness, this content must fulfil the three criteria that were established above, namely (1) non-accidental self-reference, (2) immediate action-relevance, and (3) immunity to error through misidentification.

As mentioned earlier, proponents of theories of nonconceptual content hold that there are various forms of nonconceptual self-consciousness. In the following, I will discuss the two that are considered to be the most fundamental forms of self-consciousness: visual perception and somatic proprioception.<sup>12</sup>

<sup>9</sup> We will return to this notion in the next section.

<sup>10</sup> The introduction of the distinction between knowledge-how and knowledge-that can be traced back to Gilbert Ryle (1945, 1949), according to whom knowledge-how is an ability, which is in turn a set of dispositions. In contrast, knowledge-that is a relation between a thinker and a proposition.

<sup>11</sup> But notice that my arguments will go through even if you reject the notion of nonconceptual content that I have sketched here.

<sup>12</sup> Insofar as other forms of nonconceptual self-consciousness that are discussed by Bermúdez and others are genuinely nonconceptual, my arguments will apply to these as well. Note, however, that I assume the difference between nonconceptual and conceptual forms of representation to be a gradual one, such that there will be cases where there is not a clear-cut decision as to whether a particular representation is to be considered nonconceptual or conceptual. And there are indeed some cases to be found in Bermúdez' discussion (in particular in his discussion of psychological self-awareness) where it seems unclear whether we are still dealing with a nonconceptual representation. For the sake of clarity of my argumentation, I will restrict my discussion here to cases that are clearly instances of nonconceptual representation.

Based on Gibson's (1979) theory of ecological objects, Bermúdez (1998, 2001) argues that the structure of visual perception contains information about the self in various ways: On the one hand, there are self-specifying structural invariants such as the boundedness of the visual field and the occlusion of parts of the visual field by various parts of the body. On the other hand, visual kinesthesia, that is the patterns of flow in the optic array and the relations between the variant and invariant features, enables the perceiver to experience her movement in the world. Finally, and perhaps most importantly, the perception of affordances, that is properties of objects in the environment that relate to the abilities of the perceiver, provides the subject with information not only about the objects that are being perceived, but also about the possibilities for action that these objects afford. According to Gibson, these affordances are directly perceived (as opposed to learned or inferred). For example, I immediately perceive the coffee mug in front of me to be within my reaching distance (but not the sugar pot, so that I have to ask my office mate to hand it to me), and I know that I can safely sit down on my office chair, as I immediately perceive it to be stable enough to afford this action. According to Bermúdez, the perception of affordances is one of the most fundamental ways in which self-specifying information is received.

In addition, proprioception provides the organism with information regarding the state of the body, such as its position in space, the position of limbs relative to other parts of the body, its muscular activity, and so on. It is therefore considered to be another important source of self-specifying information.<sup>13</sup>

Proponents of theories of nonconceptual content argue that both the self-specifying information in ecological perception and somatic proprioception are primitive forms of self-consciousness because (1) they deliver the organism with self-specifying information, that is, they are about the self (or they self-refer), (2) this self-specifying information is relevant for action, and (3) it is immune to error through misidentification due to the fact that it is necessarily information about the organism itself. The basic idea is that ecological perception delivers information that is necessarily about *one's own* spatio-temporal position relative to other objects in the environment and about *one's own* possibilities for interaction with these objects. Likewise, somatic proprioception delivers information that is necessarily about *one's own* body. If I perceive the mug in front of me to be within reaching distance, no question can arise as to *who* it is that can reach the mug. Likewise, if I proprioceive my legs to be crossed, there is no question as to *whose* legs are crossed—perception and proprioception do not require any self-identification. Thus, it is argued, ecological perception and somatic proprioception seem to fulfil all three criteria for self-consciousness, as identified above. They deliver information that is about oneself—about one's position relative to other objects in the environment, about the state of one's body, and so on. This information is used to guide one's intentional actions, such as grasping the cup in front of one. And the information is thought to be immune to error through misidentification, as no question can arise as

<sup>13</sup> Note that one should distinguish between proprioceptive information processing at the subpersonal level and proprioceptive awareness. Only the latter is of interest here. (Likewise, our concern throughout the article is not with subpersonal perceptual processes, but rather with perceptual representations at the conscious level.)

to whether it is *my* body that I proprioceive, or whether it is *me* that is positioned in such-and-such an orientation and distance from the computer monitor in front. Moreover, they are nonconceptual forms of representation since they do not require the exercise of conceptual abilities and can clearly be attributed to creatures that cannot be attributed with concept possession, such as animals and infants.

This line of thought seems *prima facie* plausible. But is this sufficient to establish the claim that self-consciousness, or the ability to think 'I'-thoughts, is indeed independent of concept possession and that we can characterize self-consciousness in terms of nonconceptual forms of representation? The aim of the remainder of this article is to show why this is not the case. Although, as I will explain below, theories of nonconceptual self-consciousness prove to be very helpful for explaining certain features of self-conscious thought, perception and bodily experience should not be considered to be forms of self-consciousness. Hence, or so I will argue, theories of nonconceptual self-consciousness are incomplete with regard to the task of providing an account of the ability to refer to oneself in thought independent from conceptual and linguistic abilities.

### 3 Implicit self-related information versus explicit self-representation

The aim of this section is to show why proponents of theories of nonconceptual self-consciousness fail to establish that perception and bodily experience amount to self-consciousness. I will argue that we need to distinguish between implicit self-related information and explicit self-representation, and that the latter is required for self-consciousness. Proponents of theories of nonconceptual self-consciousness fail to consider this distinction. Moreover, I will show that ecological perception and bodily experience do not amount to explicit self-representation and hence are not forms of nonconceptual self-consciousness. Thus, theories of nonconceptual self-consciousness are incomplete insofar as they only establish the existence of implicit self-related information in perception and proprioception, but not the existence of explicit self-representation.

The basic idea behind the arguments of this section is that while perception and bodily experience contain information that is in fact related to the organism itself, this information is not being represented as being about the organism itself. To put it briefly: The nonconceptual representational contents of perception and bodily experience do not constitute forms of self-consciousness because they do not contain any self-referring component; rather, they are self-less.

This is not to say that theories of nonconceptual self-consciousness do not point to something important in the context of the attempt to understand how beings come to have the ability to think 'I'-thoughts. Proponents of theories of nonconceptual self-consciousness are correct in pointing out the different kinds of self-specifying information that are implicit in perception and bodily experience. And, as we will see in Sect. 4, the presence of this information does in fact play an important role when it comes to explaining the essential features of 'I'-thoughts, in particular their immunity to error through misidentification. Nonetheless, as I will argue in the following, implicit self-related information does not amount to self-representation

and more work needs to be done if we are to understand how the ability for self-conscious thought arises out of nonconceptual ways of representation.

In order to spell out this basic idea I will present two arguments that will show where theories of nonconceptual self-consciousness fall short of what they attempt to explain. The first argument (Sect. 3.1) will appeal to the notion of the self as an ‘unarticulated constituent’ in order to show why the self-specifying information in perception and proprioception should be understood in terms of implicit self-related information rather than explicit self-representation, and why the latter is required for self-reference, and hence for self-consciousness. The second argument (Sect. 3.2) will demonstrate that the immunity principle cannot apply to nonconceptual forms of representation. So theories of nonconceptual self-consciousness can only account for one of the three essential features of self-conscious thoughts—their immediate implications for action. This is because they fail to distinguish between implicit self-related information, which is required for action-guidance, and explicit self-representation, which is required for (non-accidental) self-reference and immunity to error through misidentification.

Nonetheless, as I will argue briefly in Sect. 4, we can use the insights provided by theories of nonconceptual self-consciousness in order to explain how it is possible to have self-referring thoughts that are immune to error through misidentification. But in order to so, we need to make intelligible the transition from implicit self-related information to explicit self-representation.

### 3.1 The self as an ‘unarticulated constituent’

As we have seen in Sect. 2, perception and bodily experience provide the organism with information that is self-related, such as information about the distance between the organism and an object in its visual field, or about the position of the organism’s limbs, its position in space, about whether it is moving or not, or about the organism’s possibilities for interacting with the environment. (And this seems to fit well with the kind of view on nonconceptual content that I sketched in Sect. 2.1.) However, as I will show here, this information does not have to be explicitly represented as being *about* the organism.<sup>14</sup> Hence, while perception and bodily experience may represent instances of conscious experience, they do not constitute forms of *self*-consciousness. To put it differently: All sentient beings are subjects of experience and experience the world from their own egocentric perspective. Because of this, perceptual content necessarily contains self-related information, which the organism must make use of in order to interact with the environment. However, not all subjects of conscious experience also have explicit self-representations or think of themselves as themselves (Baker 1998). Just because

<sup>14</sup> Following Dienes and Perner (1999), I take it that a “fact is explicitly represented if there is an expression (mental or otherwise) whose meaning is just that fact; in other words, there is an internal state whose function is to indicate that fact” (Dienes and Perner 1999, p. 736). That is to say that a fact or state of affairs is represented explicitly when the mental state in question contains a component that directly refers to this fact or states of affairs. In contrast, a fact or state of affairs is implicit in a mental representation when the mental state in question does not contain a component that directly refers to this fact, but when this fact or state of affairs is conveyed as part of the contextual function of the mental state.

all conscious experience is experience *for* a subject, it is not also experience *of* a subject.

Let me illustrate this thought by means of an example<sup>15</sup>: The squirrel that is currently sitting on a branch of the walnut tree in front of my window and is about to jump to a neighboring branch has a certain perspective and is acting from this perspective. Assuming that it is acting intentionally, that is to say, assuming that we can ascribe representational content to the squirrel, one possible explanation of the squirrel's behavior may be something like the following: The squirrel sees some walnuts on the neighboring branch, it wants the nuts and consequently it jumps onto the branch. But in order to behave like this, the squirrel does not need to have any explicit representation of itself (for instance as an individual agent or perceiver, or of having a particular perspective)—it only needs to represent the tree and the nuts. To be sure, the content of the squirrel's tree-representation will be partly determined by the squirrel's abilities for interacting with the tree, by its distance from the neighboring branch and its abilities to jump, for example. That is to say that the squirrel will perceive the tree in terms of the kind of actions it affords. Nonetheless, the squirrel does not need to represent its distance to the tree and its perspective on the tree *as such*. Rather, the self-specifying information that is contained in the squirrel's perception of the walnut tree can remain implicit (as part of the function of the mode of perception, so to speak)—the information that the squirrel receives about its environment does not have to be represented to the squirrel *as* being self-related. In particular, the squirrel's perception need not contain any self-referring component. Implicit self-related information is sufficient for action guidance; it need not be represented *as* self-related in order to be utilized by the squirrel's action guiding systems.

It might seem to be the case that the squirrel's representation involves an explicit self-reference because in order to get to the nut, the squirrel has to locate the branch holding the nut relative to *its own* position; it has to know where the branch is relative to *its own* body in order to successfully perform the jump. That is, to use terminology introduced by Perry, it has to have a grasp of the agent-relative role that the branch plays in terms of the squirrel's means/end reasoning of how to get to the nut. And this, in turn, might seem to require a representation of the agent. Similarly, to use an example by Perry, when using a fax machine, “[...] I have to move my fingers a certain distance and direction from *me*. It isn't enough to know where the buttons were relative to one another, or where the fax machine was in the building or the room. I had to know where these things were relative to *me*. It seems then, that these basic methods [i.e. methods to find out about the agent-relative roles that objects in our environment play; KM] already require me to have some notion of myself.” (Perry 2000, p. 327) This intuition is what seems to be driving the idea that the types of nonconceptual content that we discussed in the previous section amount to self-consciousness—they seem to require some form of self-representation.

But this view is mistaken. Having a grasp of the agent-relative role of the objects around oneself it is not the same as having a representation of oneself. The squirrel's

<sup>15</sup> This example is inspired by Baker (1998).

representation does not need to be *about* itself, it does not need to contain a self-referring component in order to be action-guiding. Likewise, my perception of the keyboard as I type these sentences need not contain any self-referring component (it can be specified as ‘keyboard in front at such-and-such an angle’ rather than ‘keyboard in front of *me*’). This is because

Sometimes all of the facts we deal with involving a certain  $n$ -ary relation involve the same object occupying one of the argument roles. In that case, we don’t need to worry about that argument role; we don’t need to keep track of its occupant, because it never changes (Perry 2000, p. 328).

In the case of agent-relative representations the argument role of the agent itself always remains the same. All the objects that the squirrel perceives are perceived from its egocentric perspective, in-relation-to-itself. Consequently, there simply is no need for the squirrel to keep track of whose perspective is represented in its perception because the perspective always remains fixed. In other words, because there is only ever one perspective for the squirrel—its own, egocentric perspective—there is no need for the squirrel to keep track of this perspective, and hence no need to represent it explicitly. So even though the squirrel’s perception of the nut *concerns* the subject of perception, since the nut is perceived relative to the squirrel’s perspective, it is not *about* the subject.<sup>16</sup> Perceptual experiences are not about the self, they are not self-referring in the way ‘I’-thoughts are—even though they contain (implicit) self-related information. What they are about are the objects of experience, such as the branch of a tree in the squirrel’s case, or the keyboard in the human case. And the same is true of a subject’s bodily experiences. These may be characterized as being about, for instance, a ‘pain in the left shoulder’ or an ‘itch in the right foot’ rather than about a ‘pain in *my* left shoulder’ or an ‘itch in *my* right foot’.

We may also frame this, to use another term from Perry, by saying that in these cases the subject itself remains an ‘unarticulated constituent’ of its perception.<sup>17</sup> According to the theory of unarticulated constituents, “we don’t articulate the objects we are talking about, when it is obvious what they are from the context” (Perry 1998, p. 11). One of the examples that Perry uses in order to illustrate this idea is the well-known case of the Z-landers. Residents of Z-land never get any information about the weather anywhere else, they may not even be aware of the existence of other places. When they say ‘It is raining’, this statement is obviously relevant to the weather in Z-land, and they are able to act accordingly (by, for example, taking an umbrella when leaving the house), even though Z-land is not explicitly represented in the utterance. Z-land figures as an ‘unarticulated constituent’ of the utterance because in order to determine the truth conditions of

<sup>16</sup> The applicability of Perry’s distinction between thoughts that concern an object and those that are about an object to the discussion about nonconceptual self-consciousness is also discussed in Meeks (2006)—though he is rather critical of this application.

<sup>17</sup> The notion of the self as an ‘unarticulated constituent’ of perception was first brought to my attention through a conference presentation by Arnon Cahen (2006). Also see Meeks (2006) for a related discussion, which also mentions Z-land and considers the analogy to perception. Note, however, that both authors come to rather different conclusions from the one defended here.

the sentence ‘It is raining’ we need a location (in this case Z-land)—the sentence will be true iff it is indeed raining in Z-land. Yet Z-land does not figure as an explicit component in the utterance, and it need not be explicitly represented either for the statement to be understood. In contrast to us, Z-landers do not need to consider the weather in other locations when thinking or speaking about weather—as a matter of fact, they are not even able to consider other locations. Accordingly, there simply is no need for there to be a component referring to Z-land in their thoughts in order for the connection to Z-land to be secured; there is no need (nor the ability) to “keep track” of the location. Z-landers’ thoughts about weather *concern* Z-land insofar as they lead to behavior that is appropriate to the weather in Z-land (e.g., taking an umbrella when leaving the house upon thinking ‘It is raining’), but Z-land does not have to be represented for this to hold (hence their thoughts are not *about* Z-land).<sup>18</sup>

Similarly, a being that never learns to distinguish its own perspective from that of others, a being who is not aware that other perspectives on the world exist, does not need to explicitly represent its own perspective in its perception and thinking. Even for beings like us, who do have a sense of their own perspective (in contrast to that of others), this perspective does not have to be represented explicitly in our perceptual experience, for this would put an unnecessary burden on our cognitive system. Perry himself draws the analogy in the following way:

What each of us gets from perception may be regarded as information concerning ourselves, to explain connections between perception and action. *There is no need for a self-referring component of our belief, no need for an idea or representation of ourselves.* [...] The eyes that see and the torso or legs that move are parts of the same more or less integrated body. And this fact, *external to the belief*, supplies the needed coordination. The belief need only have the burden of registering differences in my environment, and not the burden of identifying the person about whose relation to the environment perception gives information with the person whose action it guides (Perry 2000, pp.182–183; italics mine).

Accordingly, contrary to what theories of nonconceptual content seem to suggest, no self-referring component of the experiential state is required in order to secure the connection between perception and action. What proponents of nonconceptual self-consciousness seem to argue is that since perception and bodily experience

<sup>18</sup> Note that in his essay ‘Thought without representation’ Perry states that “Z-land comes in not as an unarticulated constituent each Z-landish weather statement is about, but a global factor that all Z-land discourse about the weather concerns”. (2000, p. 179) This is to distinguish the Z-land case from cases where we talk about weather without articulating a location because it is obvious from the conversational context what the location is—in our case, the location does come in as an unarticulated constituent of each weather statement. In contrast to Z-landers we are aware of other places and thus have to “track” which location the conversation is about (whether or not that location is articulated). But since Z-landers are not aware of other places, no such “tracking” has to take place for them—their beliefs and assertions carry a “lesser burden” compared to ours, so to speak. Accordingly, “the Z-landers assertions and beliefs *concern* Z-land, but are not *about* Z-land” (Perry 2000, p. 179: emphasis original). What I am claiming here is that the case of perception and proprioception is analogous to the case of weather-reports in Z-land.

provide information that is relevant for guiding the intentional actions of an organism, their content must be about the self, in other words, it must refer to the self. But as we have just seen, this is not so. Perception and bodily experience can guide an organism's intentional behavior in virtue of containing agent-relative information (due to the way in which perception and proprioception function) without being *about* the organism, or containing a self-referring component. There simply is no need to represent the subject of perception and bodily experience, as it always remains the same. (Hence it can remain 'unarticulated'.)

Notice, though, that Perry is concerned with the (conceptual) content of belief and its truth conditions—the notion of an 'unarticulated constituent' refers to an element that figures in the determination of the truth conditions of a belief without being separately represented. In contrast, what I am concerned with here is the (nonconceptual) content of perception (and proprioception) itself, not the content of the beliefs that a subject can form based on her perceptions. I nevertheless take it that we can use Perry's insights in order to elucidate how we should specify the content of the former. On the view that I propose here, perception and proprioception provide the organism with information that is (necessarily) self-related and that enables the organism to coordinate its perception and action in order to engage and interact with its environment. But precisely because the perceptual content is always related to one particular subject, the subject itself does not need to be identified or represented in the perceptual content—similar to the linguistic case of the *Z*-landers in which the location does not have to be represented because it always remains fixed. That is to say that the point I want to make is that the contents of perception and proprioception themselves do not contain a self-referring component (independent as to whether a belief that is formed based on the former does).

Notice also that, on the view I have sketched and endorsed in Sect. 2.1, nonconceptual content should not be characterized in terms of truth conditions. This raises the question as to how the notion of an 'unarticulated constituent' can be transferred from belief to perception. As I have suggested, nonconceptual content provides the subject with a form of knowing-how (e.g., 'knowing-how-to-interact-with-the-tree-in-front'), not with a form of knowing-that (e.g., 'that there is a tree in front'). While the latter is truth-evaluable, the former is not. However, the subject could, if she possessed the relevant concepts, form a belief on the basis of her perception (such as the belief 'that there is a tree'), which would count as a true perceptual belief if and only if the subject of the perception does indeed stand in a perceptual relation to the tree. That is to say that the subject does play a role in determining the truth conditions of a belief that is formed based on her perception.<sup>19</sup> But this is not to say that the subject (or the perceptual relation between the subject and the tree) must be represented in perception (nor in the content of the resulting perceptual belief). Likewise, in the case of proprioception, a belief that 'legs are crossed' which is formed on the basis of the proprioception of one's crossed legs will be true if the property in question (that of crossed legs) is in fact instantiated in

<sup>19</sup> In Recanati's (2007, 2009) terminology, the subject is part of the circumstance of evaluation of the content of the resulting perceptual belief.

oneself. But that is not to say that either the self or the instantiation relation must be represented in proprioception (cf. Recanati 2009, pp. 266–267). Thus, we can say that the self is an ‘unarticulated constituent’ of the content of perception and proprioception in virtue of the fact that the beliefs that are formed *on the basis of this content* will have to be truth-evaluated with respect to the subject, even though the subject herself (and the perceptual or instantiation relation) are not represented in perception and proprioception (and neither do they have to be represented in the corresponding belief—though on occasion they might be). So the notion of the self as an ‘unarticulated constituent’ of perception and proprioception is derivative upon the role that the self plays in evaluating the relevant beliefs that the subject could form on the basis of their perceptual and proprioceptive states.

Finally, let me point out that a being who possesses the first-person concept could, of course, on the basis of having a perceptual experience of a tree in front, not only form the belief that there is a tree, but could also form the explicit first-person judgment ‘I see a tree’, thereby explicitly self-ascribing this experience. (In fact, a subject who possesses the first person concept not only could, but would also be justified in forming this judgment.) That is to say that the subject *can*, but does not have to, be represented in the content of a (conceptual) belief that is formed based on the (nonconceptual) content of her perception or proprioception. But the fact that perceptual content can, in the case of beings who possess the first person concept, provide a subject with a basis for forming this explicit judgment does neither imply nor require that the perceptual content itself must contain an explicitly self-referring component. And it is only the latter that we are concerned with here.

What this discussion shows is that one can have self-related information, which is sufficient for action-guidance, without any self-representation. It also shows that insofar as the self should be understood as an ‘unarticulated constituent’ of perception or proprioception, the latter do not self-refer (they are not about the self). Yet (non-accidental) self-reference was supposed to be a necessary condition for self-consciousness. Recall our analysis from Sect. 1. Self-consciousness was defined as the ability to think ‘I’-thoughts, where ‘I’ thoughts are such that they non-accidentally refer to the subject entertaining them. That is to say that in having an ‘I’-thought the subject cannot but be aware of the fact that the thought is about themselves. Since the contents of perception and proprioception neither contain a self-referring component, nor are about the self, perception and bodily experience cannot be considered to be forms of self-consciousness. In conclusion, theories of nonconceptual self-consciousness are incomplete. To be sure, they can account for the action-guiding nature of certain representational states, which is one important characteristic of ‘I’-thoughts. But they fail to establish that the representational states in question actually represent the self (rather than containing implicit self-related or agent-relative information).<sup>20</sup>

<sup>20</sup> You may have noticed that this argument is somewhat independent of the question as to whether the content of perception and proprioception is indeed best characterized as being nonconceptual. Although there are good reasons to think that this is the case, even on a conceptualist theory of perception and proprioception the same point still applies. For whether or not the content of perception and proprioception is conceptual or nonconceptual, the fact remains that no explicit self-representation is required in order to account for the action-guiding function of this content.

Notice that there are also phenomenological reasons for resisting the thought that every conscious experience is also self-conscious. When I am engrossed in a perceptual experience of a picture in front of me or focused on an activity, such as reading a book or writing an article, I am not also aware of myself. I am focused on the picture I am looking at, or on the words on the paper and their meaning, but not on myself—there simply does not seem to be an element in my experience that corresponds to the self. In fact, if I were self-conscious while enjoying my perception or performing an activity this self-awareness would rather be a distraction from the task at hand.<sup>21</sup> So although conscious experience is always experience for a subject and is experienced from the perspective of the experiencing subject, this is not to say that the subject must also be aware of itself or of having this particular perspective.

### 3.1.1 Potential objections

One might object that contrary to the above arguments, the self does directly appear as an object in visual perception and somatic proprioception. For instance, with regard to the content of perceptual experience, Bermúdez argues that the self “has a place in the content of perceptual experience in virtue of the self-specifying information that is an integral part of that perceptual experience” (1998, p. 108). As we have seen above, this information consists in (a) information about bodily invariants that bound the field of vision, (b) information from visual kinesthesia about the movement of the perceiver, and (c) information about the possibilities for action and reaction that the environment affords the perceiver (cf. Bermúdez 1998, p. 114). However, on my view, the structural invariants that make up the self-specifying information in question are part of what constitutes the egocentric perspective of the subject. It is true that (visual) perception is necessarily perspectival<sup>22</sup>—in perceiving the world, the subject necessarily does so from its own egocentric perspective. However, it is one thing to perceive the world from an egocentric perspective, and quite another to be aware of one’s perspective in first-person terms. In other words, to represent a bounded visual field is not to represent the bounds of the visual field.<sup>23</sup> And as I have argued above, there simply is no need

<sup>21</sup> See Schear (2009) for a relevant discussion of Sartre’s phenomenological analysis of the experience of being immersed in reading a book. And see Goldberg et al. (2006) for a neuroscientific study that found a segregation between the patterns of activity associated with a demanding perceptual task and those associated with self-reflection. Moreover, the study showed that regions that showed enhanced activity during self-reflection were inhibited during the demanding perceptual task.

<sup>22</sup> It is an interesting question—that I will not address here—whether this is true also for other types of perceptual experience, such as smell.

<sup>23</sup> The objection seems to commit the error of confusing the vehicle with the content of representation (Millikan 1993). According to Millikan, one mistake often committed is to transfer properties of the vehicle of representation to the content, and vice versa. One example “involves importing and exporting the limit of a representation. If there is a limit to what is represented, there is a corresponding limit to the vehicle of representation, and a limit to the vehicle of representation is then exported to be a representation of the limits of the represented. But the limit of a represented content is not a representation of a limit.” This mistake seems to be in play when it is concluded that because the visual field is bounded, the bounds of the visual field are represented in the visual field.

for the perspective itself to be represented in perception. Neither is it necessary for the subject to represent information about the possibilities for interaction afforded by the environment as being *about* the self. As we have seen above, having a grasp of the agent-relative role of the objects around oneself it is not the same as having a representation of oneself as an agent.

A related objection might be that, since perception always contains information about the world as well as about the perceiving organism, there has to be a mechanism that is able to distinguish between self and non-self/world (Vosgerau 2009).<sup>24</sup> As we have seen in the previous section, every perceptual input carries information about both the perceiving subject and the world. Accordingly, the organism must somehow divide the input into self-related and world-related information. Given the need for this computational differentiation, one might argue, the self is, after all, part of the explicit representational content of perception. But this would be question-begging, since the point of the argument above was precisely that *despite* the fact that perception and proprioception involve implicit self-related information, the self is *not* explicitly represented. The mechanism involved in the self-world-distinction might be able to explain how the implicit self-related information is processed at a subpersonal level, but this does not imply that the system represents the self at the personal level. Personal or animal level explanations aim at representational content that is accessible to the subject in question such that it can provide the basis for the kind of practical reasoning processes associated with intentional action. This is the level we are concerned with when we are considering the possibility of self-conscious thought. Subpersonal level explanations on the other hand describe processes that might constrain personal level representations in interesting ways, but that are not accessible to the organism as such. Indeed, the explanations that have been suggested for such a self-world-distinction mechanism all seem to locate this mechanism at the subpersonal level.<sup>25</sup>

These subpersonal mechanisms enable an organism to perceive and interact with the world; they provide a subpersonal level explanation for the fact that perception contains (implicit) self-related information. But again, it is one thing to claim that for an organism to be able to interact with its environment the content of perception must be self-related (which is a claim that I do not object to) and to specify a subpersonal mechanism for this, but it is another thing to claim that this self-related information must be represented explicitly at the personal level as being about the self. It is certainly the case that subpersonal level explanations can contribute—as enabling or making-possible mechanisms—to an understanding of personal-level phenomena (Hurley 2008; McDowell 1994). For instance, breakdowns in efference copy mechanisms have been suggested as part of the explanation for the phenomenon of thought insertion in schizophrenia (Frith 1992). Nonetheless, we cannot simply project properties between the personal and subpersonal levels, as

<sup>24</sup> This objection was proposed to me by Gottfried Vosgerau during a conference presentation.

<sup>25</sup> These explanations usually involve the comparator model, which is a subpersonal, unconscious mechanism that ensures motor control by comparing a copy of the efference signal of a motor command with the afference (including reafference) signal (von Holst and Mittelstaedt 1950).

this would be confusing the vehicle with the content of representation (Millikan 1993).

So, although it makes sense to postulate a self/world distinction mechanism at the subpersonal level, the existence of such a mechanism does not undermine the argument presented above. Quite to the contrary, the argument that the self does not need to be explicitly represented in perception and bodily experience was based precisely on the idea that perception and bodily experience necessarily involve self-related information. It is precisely because of the presence of this information, because of the fact that the argument role of the self remains fixed, that the self does not have to be explicitly represented in order for perception and bodily experience to be action-guiding. So rather than posing a problem, the postulation of a self-world-distinction mechanism is quite compatible with the notion of the self as an ‘unarticulated constituent’ because it can explain the subpersonal mechanism in virtue of which perception and bodily experience are self-related.

The burden of proof is on the proponent of a theory of nonconceptual self-consciousness to show that contrary to my argument, perception and proprioception must contain explicit self-representations, rather than merely implicit self-related information, in order to enable the organism’s interaction with the environment.

However, one might still wonder whether the arguments presented equally apply to perception and bodily experience. After all, while it seems plausible that perception really is just about the external objects of perception, proprioception seems to be about one’s own body. And, given a basic materialist sentiment, one might think that since one’s body is oneself, proprioception therefore really is about oneself, in the sense that the self appears directly as an object of proprioception. As Bermúdez puts it: “What somatic proprioception offers is an awareness of the body as a spatially extended and bounded physical *object* that is distinctive in being responsive to the will.” (1998, p. 150; emphasis mine.) According to Bermúdez, the body literally appears as an object of proprioception, as the latter provides the subject with an awareness of the limits of the body.

However, the representation of bodily sensations differs from the perception of objects in the outside world in an important respect. For while perception locates objects relative to an egocentric frame of reference, bodily sensations do not locate its objects of experience relative to an egocentric frame of reference. Rather, they provide this frame of reference in the first place (Gallagher 2003). Moreover, while it is true that proprioception provides the organism with information about its body, and that this information necessarily is about the organism’s *own* body, and thus about the organism, it is not being represented *as* such. That is to say it is not being represented *as* being about the own body (rather than anybody else’s). In other words, because proprioception is necessarily of a unique subject, the subject itself does not have to be represented explicitly. The argument role of the agent in a subject’s proprioceptive representational states always remains the same and hence does not have to be explicitly represented. The constitutive relation to the subject is provided by the functional role of the proprioceptive state (which is such that it necessarily represents bodily states and properties of one’s own body), not by its explicit representational content. So if one were to specify the content of such a representation, for instance of one’s legs being crossed or of the pain in one’s

shoulder, it would be ‘legs are crossed’ rather than ‘*my* legs are crossed’ and ‘pain in left shoulder’ rather ‘pain in *my* left shoulder’, respectively.<sup>26</sup>

In sum, perception provides us with the experience of objects in our environment and proprioception provides us with bodily sensations, but the fact that these are necessarily experiences for a subject obviates the need for an explicit representation of the subject.

### 3.2 Nonconceptual content and immunity to error through misidentification

Having shown how representational content can implicitly contain self-related information without explicitly representing the self, I will now go on to argue that, for structural reasons, nonconceptual content cannot support explicit self-representations.

As we have seen above, nonconceptual content is defined negatively, in opposition to conceptual content. While conceptual content is generally considered to meet Evans’ (1982) Generality Constraint, nonconceptual content is not. One might take this to mean that nonconceptual content is noncompositional (Heck 2007; Meeks 2006; Toribio 2008). If that is so, this would mean that nonconceptual content does not have the necessary structural features required for an explicit predication relation of the form ‘*a* is *F*’ (which could be decomposed into object and predicate) or, for that matter, ‘*I*  $\varphi$ ’, and hence does not have the necessary structural features required for an explicit self-representation (as in “I am hungry”, “I am in pain”, etc.). The structure required for explicit self-representation can only come with content that is compositional, and thus can only come with conceptual content. So far, this argument lends further support to the claim established in the previous section, namely that there is no explicit self-representation in perception and proprioception, although the argument of this section is based on the noncompositionality of nonconceptual content and not on the fact that self-representation is not necessary for perception and action guidance.

However, the fact that nonconceptual content is thought to be noncompositional also implies that nonconceptual content cannot be said to be immune to error through misidentification. Let us recall the basic definition of immunity to error through misidentification. As we have seen in Sect. 1, a thought that is immune to error through misidentification allows for the possibility that it misrepresents the property that is being ascribed, while it cannot misrepresent the subject purportedly possessing that property. What this means, though, is that the notion of immunity is applicable only to judgments, and not to nonconceptual representations.<sup>27</sup> In other words, the notion of a representational distinction between subject and property, which would allow for a misrepresentation of the property that is being ascribed

<sup>26</sup> Which is not to say that the subject could not, on the basis of her bodily experience, form the explicit belief that *her* legs are crossed (by applying the first-person concept). But, as I argued above, the fact that the content of her experience provides her with sufficient warrant to form this belief (due to the fact that her experience is necessarily self-related) does not imply (or require) that the experiential content itself contains an explicitly self-referring component. Rather, the subject becomes self-aware only via the explicit self-ascription of the experience by means of applying the first person concept.

<sup>27</sup> I am indebted to Arnon Cahen for helpful discussions of this point.

while leaving the representation of the subject intact, is a matter in the realm of conceptual contents. Nonconceptual content does not represent the subject of a thought or perception independent from the property that is being ascribed. Rather, the representational content is in a format in which the property is necessarily implicitly self-related, and in which the subject is not represented at all. (In other words, the self-relatedness of the representation is provided by a fact that is external to the representation.) What this implies is that the question as to whether nonconceptual content is immune to error through misidentification cannot even arise. In other words: the question as to whether nonconceptual content can be immune to error through misidentification can only arise at the level of *judgment*, not at the level of nonconceptual content. It is a category mistake, so to speak, to try to apply the notion of immunity at the level of nonconceptual content.<sup>28</sup>

This is not to say that we can in fact be mistaken as who the subject of our perceptual or proprioceptive states is. The fact remains that ecological perception and bodily experience contain information that is self-specifying, and that this is necessarily so. Rather, because this information is represented in a way in which the subject of perception is not a component of the explicit representational content (because the representational content is implicitly self-related), we cannot say that there is a possibility of predicate-misrepresentation while leaving the representational subject intact—both because nonconceptual content does not have the structure that would be required for a representation of a subject and a predicate, and, more importantly, because there simply is no representational subject that could be represented independently from the representational predicate. But the possibility of explicitly representing the subject independent from the predicates that are ascribed to the latter—at least in principle—is a necessary precondition for the notion of immunity to gain any traction. Therefore, ecological perception and bodily experience can at best be said to provide the *basis* for the formation of first person judgments that can have the property of immunity and thus represent paradigmatic forms of self-consciousness, but they do not by themselves represent such paradigmatic forms of self-consciousness.

### 3.2.1 Potential objections

It might be objected that nonconceptual content could fail to satisfy Evans' Generality Constraint without being entirely noncompositional. Indeed, Peacocke (1992, Ch. 3) argues explicitly in favor of what he calls proto-propositional content at the nonconceptual level, and Bermúdez (1998, Ch. 3) likewise argues that nonconceptual content is compositional, and that we can ascribe “proto-beliefs” to

<sup>28</sup> Notice that the argument presented here is different from the one in Meeks (2006), although it was inspired by the latter. Meeks argues that due to the noncompositionality of nonconceptual content, it cannot accurately represent the subject of a self-ascription while misrepresenting the property—it can only misrepresent *tout court* (or else fail to count as genuine representational content). Hence, according to Meeks, nonconceptual content cannot be immune to error through misidentification. In contrast, what I am arguing is that the notion of immunity to error through misidentification cannot apply to nonconceptual content in the first place, so nonconceptual can neither be said to be immune, nor can it be said not to be immune.

nonconceptual creatures. Thus, one might accept something like weak compositionality at the nonconceptual level, and this, in turn, might be sufficient for the notion of immunity to gain traction.

But at this point one surely must ask whether there really is a difference in kind between conceptual and nonconceptual content (cf. Meeks 2006). If the weak compositionality at the nonconceptual level is supposed to support immunity to error through misidentification, then it would have to allow for the representation of a subject-component as well as for the representation of a predicate-component, and moreover it would have to allow for these components to be represented independently from each other (such that the latter could be misrepresented while leaving the representation of the former intact). It is hard to see why this would not constitute a case of concept possession.

So what are we to make of the distinction between conceptual and nonconceptual content on a view that wants to maintain the possibility of weak compositionality (albeit strong enough to allow for immunity) at the nonconceptual level? Bermúdez argues that while creatures that possess nonconceptual content are sensitive to the truth of inferential transitions, in order for a creature to be credited with concept possession, it must also be able to identify and provide reasons for the inferences it is making. The latter is a paradigmatically linguistic ability, and therefore, there cannot be conceptual representations in the absence of linguistic abilities:

Providing justifications is a matter of identifying and articulating the reasons for a given classification, inference, or judgment. It is because prelinguistic creatures are in principle incapable of providing such justifications that the priority thesis is true. Mere sensitivity to the truth of inferential transitions involving a given concept is not enough for possession of that concept. Rational sensitivity is required, and rational sensitivity comes only with language mastery (Bermudez 1998, p. 71).

Thus, for Bermúdez, concept possession requires not only the ability to draw inferences, but also the ability to identify and articulate reasons for one's inferential transitions. On such a view we might be able to ascribe a level of compositionality to nonconceptual contents that does allow for subject–predicate ascriptions without blurring the distinction between conceptual and nonconceptual content. But this very demanding view on concept possession comes at a high cost. As Meeks (2006) has pointed out, imposing such a requirement on concept possession may prevent us from ascribing concepts where we would otherwise be comfortable doing so, including in cases which Bermúdez himself cites as examples of concept possession. So it seems that if one wants to grant compositionality to nonconceptual contents, such that the latter can support representations that are immune to error through misidentification, the distinction between conceptual and nonconceptual content loses its force, unless one imposes implausibly strict criteria on concept possession.

On my view, rather than modeling nonconceptual content along the lines of conceptual content, it is more plausible, and explanatorily more forceful, to maintain a distinction in kind between conceptual and nonconceptual contents, which, as I suggested in Sect. 2, locates the latter in the realm of knowledge-how, rather than in the realm of the identification of objects.

However, even if we were to allow for a certain level of compositionality of nonconceptual representations, this would not get around the problem that was discussed in the Sect. 3.1, namely the fact that the self is not represented in perception and proprioception, which is to say that the latter do not contain a self-referring component. (For no self-referring component is needed to account for the essential self-relatedness and action guiding function of perception and bodily experience.) That is to say that the content of perception and proprioception does not contain a component that could stand for the ‘I’ of a self-conscious thought. And since the immunity principle can only apply to representational states that contain a self-referring component, such that this component can be represented independently from the predicate component, the immunity principle cannot apply to the content of perception and proprioception. So the immunity principle would still fail to apply to states with nonconceptual content, even if we did allow for the latter to be compositional.

The upshot of this section is that immunity to error through misidentification is a property of (conceptual) judgments, not of the (nonconceptual) perceptive or proprioceptive states based upon which such judgments might be made. The notion of immunity cannot apply to states with nonconceptual content, as the content of these states is noncompositional and hence cannot represent subject and predicate independent from each other. Moreover, and more importantly, as I have argued in the previous section, perception and bodily experience do not contain a self-referring component, despite the fact that their content is necessarily self-related. Therefore, states with nonconceptual content cannot represent forms of genuine self-consciousness for they fail to fall under the category of representations that can be said to be immune to error through misidentification. They can at best provide the *basis* for first personal judgments that are immune.

#### 4 Immunity reconsidered

But how can representational states with nonconceptual content provide the basis for first-personal judgments that are immune to error through misidentification, if they are not to be considered immune themselves (and, moreover, if they do not contain a self-referring component)? In order to better understand this, again, a turn to philosophy of language will prove helpful. According to Recanati (2007, 2009), it is exactly those mental states that necessarily involve implicit self-related information but no explicit self-reference at the level of content that underwrite first-person judgments that are immune to error through misidentification. He calls these mental states ‘implicit *de se* thoughts’, which is somewhat misleading, as, according to the argument I presented above, it is hard to see how we can make sense of a *de se* (or ‘I-’) thought that involves no self-reference. These terminological quibbles notwithstanding, Recanati’s work can help us shed light on the phenomenon of immunity to error through misidentification. Recall that, as we have seen earlier, perception and bodily experience necessarily contain information about the experiencing subject, but they do so *not* because the subject is part of the explicit content of the experience, but rather because the experience is gained in such a way

that it is necessarily related to the subject. For instance, in seeing a tree in front of me, I cannot gain information about anybody else's visual experience but my own. Similarly, I cannot, through proprioception, gain information about anybody else's body but my own. Now, according to Recanati, if "on the basis of his experience, the subject forms the explicit first person judgment that his own legs are crossed (in contrast, say, to his neighbour's), the judgment will retain the immunity to error through misidentification which characterized the proprioceptive experience on which it is based" (2009, p. 12). This is so because the explicit judgment does not require any additional evidence or extra premise (and in particular no identity premise). Now, again, in the light of the arguments discussed in this paper, Recanati's formulation is misleading for, strictly speaking, the notion of immunity does not apply to the (nonconceptual) representational content of proprioceptive experience. For neither is the self explicitly represented (so there is no subject in the representational content), nor does the representational content have the adequate structure to account for a representation of the subject independent from the predicate. This misleading formulation notwithstanding, the important point that Recanati makes is that an explicit first-person judgment can be immune to error through misidentification insofar as no additional evidence base (i.e., no extra premise) is introduced when the judgment is formed based on an experience that is necessarily self-related (if not self-representational). Hence, there is no additional source of error (and in particular no source of misidentification).

So proponents of theories of nonconceptual self-consciousness do make a very important point when they refer to the essential self-relatedness of certain types of perceptual experience, even if they conflate implicit self-related information and explicit self-representation. For it is this essential self-relatedness of certain ways of gaining knowledge about the world that ultimately underwrites the immunity to error through misidentification of certain 'I'-thoughts.

Now obviously, this still leaves us with the question of what exactly explicit self-representation amounts to and how we are to make intelligible the transition from implicit self-related information to explicit self-representation. This is not the place to answer this question, but I believe that to understand this transition we need to give an account of the acquisition and application of the first person concept. Moreover, I think that this transition involves an awareness of other minds, as the explicit representation of one's own mental and bodily states requires the ability to contrast these with the mental and bodily states of other subjects. That is to say that I take the notion of self-consciousness to be a contrastive notion.<sup>29</sup>

## 5 Conclusion

We can conclude that the nonconceptual representational contents of perception and bodily experience are neither self-representational, nor do they fall under the category of representations that can be said to be immune to error through misidentification. Proponents of theories of nonconceptual self-consciousness point

<sup>29</sup> For further discussion see Musholt ([forthcoming](#)).

to the essential self-relatedness of certain forms of nonconceptual representation (in particular perception and bodily experience) and to their relevance when it comes to explaining the behavior of the subjects entertaining them. While these points are correct and important, they do not suffice to show that perception and bodily experience represent the self. In other words, while theories of nonconceptual self-consciousness demonstrate that nonconceptual representations, such as the content of ecological perception and proprioception, are implicitly self-related, they fail to establish that they are explicitly self-referring. And it is the latter that is required for an account of self-consciousness (in the sense of having the ability to think ‘I’-thoughts). Thus, in order to give a complete account of the genetic and epistemic sources of self-consciousness, we will need to make sense of the transition from representations that are implicitly self-related to explicit self-representations.

**Acknowledgements** Earlier versions of this paper were presented at the Cognitive Science Symposium at the CUNY Graduate Center in New York, the Philosophy of Mind Colloquium at the Humboldt University of Berlin, the ESPP2009 conference in Budapest, and the LSE. I am grateful to members of the audiences at these meetings for helpful comments and discussions. I would also like to thank Alex Byrne, Arnon Cahen, Tomis Kapitan, Michael Pauen, Francois Recanati, David Rosenthal, Patrick Wilken and an anonymous referee for this journal for instructive discussions and/or helpful comments on previous drafts of this paper.

## References

- Baker, L. R. (1998). The first-person perspective: A test for naturalism. *American Philosophical Quarterly*, 35(4), 327–348.
- Bermúdez, J. (1998). *The paradox of self-consciousness*. Cambridge: MIT Press.
- Bermúdez, J. L. (2001). Nonconceptual self-consciousness and cognitive science. *Synthese*, 129(1), 129–149.
- Bermúdez, J. L. (2002). The sources of self-consciousness. *Proceedings of the Aristotelian Society*, 102, 87–107.
- Cahen, A. (2006). *The implicit self in perception*. Oxford: ASSC10.
- Castaneda, H. N. (1966). He: A study in the logic of self-consciousness. *Ratio*, 8(2), 130–157.
- Castaneda, H. N., Hart, J. G., & Kapitan, T. (1999). *The phenomeno-logic of the I: Essays on self-consciousness*. Bloomington: Indiana University Press.
- Chisholm, R. M. (1981). *The first person, an essay on reference and intentionality*. Minneapolis: University of Minnesota Press.
- Cussins, A. (1990). Content, conceptual content, and nonconceptual content. In Y. H. Gunther (Ed.), *Essays on nonconceptual content* (pp. 133–163). Cambridge: MIT Press.
- Dienes, Z., & Perner, J. (1999). A theory of implicit and explicit knowledge. *Behavioral and Brain Sciences*, 22(05), 735–808.
- Dummett, M. A. (1973). *Frege: Philosophy of language*. Cambridge: Cambridge University Press.
- Evans, G. (1982). *The varieties of reference*. Oxford: Oxford University Press.
- Frank, M. (2007). Non-objectal subjectivity. *Journal of Consciousness Studies*, 14, 152–173.
- Frith, C. D. (1992). *The cognitive neuropsychology of schizophrenia*. Hove: Psychology Press.
- Gallagher, S. (2003). Bodily self-awareness and object-perception. *Theoria et Historia Scientiarum: International Journal for Interdisciplinary Studies*, 7, 53–68.
- Gibson, J. J. (1979). *The ecological approach to perception*. Boston: Houghton Mifflin.
- Goldberg, I. I., Harel, M., & Malach, R. (2006). When the brain loses its self: prefrontal inactivation during sensorimotor processing. *Neuron*, 50(2), 329–339.
- Heck, R. G. (2000). Nonconceptual content and the “Space of reasons”. *The Philosophical Review*, 109(4), 483–523.
- Heck, R. (2007). Are there different kinds of content? In B. McLaughlin & J. Cohen (Eds.), *Contemporary debates in philosophy of mind* (pp. 117–138). Oxford: Blackwell Publishing Ltd.

- Hurley, S. (2006). Making sense of animals. In S. Hurley & M. Nudds (Eds.), *Rational animals?* Oxford University Press.
- Hurley, S. (2008). The shared circuits model (SCM): How control, mirroring, and simulation can enable imitation, deliberation, and mindreading. *Behavioral and Brain Sciences*, 31(01), 1–22.
- Hurley, S. L. (1997). Nonconceptual self-consciousness and agency: Perspective and access. *Communication and Cognition*, 30(3/4), 207–248 (Part 1 of special issue: Approaching consciousness).
- Legrand, D. (2007). Pre-reflective self-as-subject from experiential and empirical perspectives. *Consciousness and Cognition*, 16(3), 583–599.
- McDowell, J. (1994). *Mind and world*. Cambridge: Harvard University Press.
- Meeks, R. (2006). Why nonconceptual content cannot be immune to error through misidentification. *European Review of Philosophy*, 6, 83–102.
- Millikan, R. G. (1993). Content and vehicle. In N. Eilan, R. McCarthy, & B. Brewer (Eds.), *Spatial representation* (pp. 256–268). Oxford: Blackwell.
- Musholt, K. (forthcoming). Self-consciousness and intersubjectivity. *Grazer Philosophische Studien*.
- Peacocke, C. (2002). Postscript: The relations between conceptual and nonconceptual content. In Y. H. Gunther (Ed.), *Essays on nonconceptual content* (pp. 318–322). Cambridge: MIT Press.
- Perry, J. (1998). Indexicals, contexts and unarticulated constituents. In *Proceedings of the 1995 CSLI-Amsterdam Logic, Language and Computation Conference* (pp. 1–16). Stanford: CSLI.
- Perry, J. (2000). *The problem of the essential indexical: and other essays*. Stanford: CSLI Publications.
- Pettit, P. (2003). Looks as powers. *Philosophical Issues*, 13(1), 221–252.
- Poellner, P. (2003). Non-conceptual content, experience and the self. *Journal of Consciousness Studies*, 10(2), 32–57.
- Recanati, F. (2007). *Perspectival thought: A plea for (moderate) relativism*. Oxford: Oxford University Press.
- Recanati, F. (2009). De re and de se. *Dialectica*, 63(3), 249–269.
- Roskies, A. L. (2008). A new argument for nonconceptual content. *Philosophy and Phenomenological Research*, 76(3), 633.
- Ryle, G. (1945). Knowing how and knowing that: The presidential address. *Proceedings of the Aristotelian Society*, 46, 1–16.
- Ryle, G. (1949). *The concept of mind*. London: Hutchinson.
- Scheur, J. K. (2009). Experience and self-consciousness. *Philosophical Studies*, 144(1), 95–105.
- Shoemaker, S. (1968). Self-reference and self-awareness. *Journal of Philosophy*, 65(19), 555–567.
- Toribio, J. (2008). State versus content: The unfair trial of perceptual nonconceptualism. *Erkenntnis*, 69(3), 351–361.
- von Holst, E., & Mittelstaedt, H. (1950). Das Reafferenzprinzip. *Naturwissenschaften*, 37, 464–476.
- Vosgerau, G. (2009). *Mental representation and self-consciousness: From basic self-representation to self-related cognition*. Paderborn: Mentis.
- Ward, D., Roberts, T., & Clark, A. (2010). Knowing what we can do: Actions, intentions, and the construction of phenomenal experience. *Synthese*, 181, 375–394.
- Wittgenstein, L. (1958). *The blue and brown books*. Oxford: Blackwell.
- Zahavi, D. (2005). *Subjectivity and selfhood: Investigating the first-person perspective*. Cambridge: MIT Press.