

Mapping the Minds of Others¹

Alexandria Boyle

Note: This is a preprint of a paper forthcoming in the Review of Philosophy and Psychology. Any citations should be to the published version.

Abstract: Mindreaders can ascribe representational states to others. Some can ascribe representational states – states with semantic properties like accuracy-aptness. I argue that within this group of mindreaders, there is substantial room for variation – since mindreaders might differ with respect to the representational *format* they take representational states to have. Given that formats differ in their formal features and expressive power, the format one takes mental states to have will significantly affect the range of mental state attributions one can make, and the ease or difficulty with which one can make them. I illustrate this by considering what it would be to take mental states to be map-like in format, showing that this would result in a distinctively limited form of mindreading. I close by articulating the significance of this for the interpretation of evidence about great ape mindreading.

1. Introduction

An emerging, though not universal, consensus in comparative psychology has it that at least some non-human animals ‘read minds’. That is to say, they ascribe mental states to others, and predict behaviour on the basis of these mental state ascriptions. There is evidence that many animals reason in this way about others’ goals and intentions, as well as their epistemic states (Bugnyar, Reber, & Buckner, 2016; Call & Tomasello, 2008; Clayton, Dally, & Emery, 2007; Hare, Call, Agnetta, & Tomasello, 2000; Hare, Call, & Tomasello, 2001; Marticorena, Ruiz, Mukerji, Goddu, & Santos, 2011). Until recently, it was thought that animals² were characteristically limited in their ability to represent epistemic states: they could represent only factive epistemic states, like perception and knowledge, and were consequently unable to pass the ‘false belief test’ (Call & Tomasello, 2008). On this basis, it’s been suggested that they understand mental states as relations holding between individuals and objects or situations, but not as representations with semantic properties like being apt for accuracy and inaccuracy. Recently, though, it appears great apes have cleared even that hurdle: their anticipatory looking behaviour suggests that they predict how others will act when they have a false belief about an object’s location (Krupenye, Kano, Hirata, Call, & Tomasello, 2016).³

¹ Acknowledgement removed for anonymous review.

² Here and throughout, I use ‘animals’ as shorthand for ‘non-human animals’; similarly, where I use ‘great apes’ I mean this to exclude humans.

³ See also (Buttelmann et al., 2017; Krachun, Carpenter, Call, & Tomasello, 2009)

Suppose we take this result at face-value, as showing that great apes represent representational states as such – as the kind of thing that can be accurate or inaccurate. Does this show that great ape mindreading abilities differ from our own in degree only, and not in kind?⁴ In this paper, I argue that it does not: there may be substantial differences between great ape and human mindreading, even if great apes represent representational states as such. One underexplored dimension along which forms of mindreading may vary concerns the representational format a mindreader takes mental states to have. Humans, I suggest, take representational states to have a linguistic representational format. If mental states were taken to have a non-linguistic representational format, I argue, the result would be a distinctive and limited form of mindreading.

As a case in point, I discuss the possibility of ‘mindmapping’ – that is, taking mental representations to be map-like. I show that this would provide an effective way of representing others’ beliefs about the spatial arrangement of their environment, but that it would otherwise impose substantial constraints on the range of belief attributions one could make. In general, the representational format one takes mental states to have will significantly affect one’s mindreading abilities – since representational formats differ in their formal features and expressive power. I close by articulating the significance of this for the study of great ape mindreading. I argue that our evidence currently underdetermines what format apes take representational states to have: consistently with our evidence, they might be mindmappers. Further empirical work might reveal whether or not this is so, thereby providing a more complete picture of great apes’ understanding of the mind.

2. Varieties of Mindreading

To be a mindreader is to be able to ascribe mental states to others. In adult humans, this often takes the form of ‘propositional attitude mindreading’.⁵ That is to say, we ascribe propositional attitudes like ‘he thinks the ice is melting’, ‘she wants to win the tournament’ and so on. We take propositional attitudes to have contents which are truth- or accuracy-apt, which can be expressed using sentences, and which stand in inferential relations mirrored by the inferential relations between the sentences used to express them. It is in virtue of these inferential relations that propositional attitudes contribute to reasoning and give rise to behaviour in the

⁴ This is suggested by the editor’s summary of the Krupenye et al. study.

⁵ I say ‘often’ here because the role of propositional attitude mindreading in human cognition has probably been greatly exaggerated (see Andrews, 2012; Bermúdez, 2003a).

way that they do. We appreciate this, and are able to predict and explain behaviour accordingly.⁶

But there might be other ways to represent the minds of others, besides representing propositional attitudes as such. In fact, standard characterisations of mindreading leave open the possibility that there might be different varieties of mindreading. One characterisation of mindreading treats it as the representation of causally efficacious unobservable states (Penn & Povinelli, 2007; Whiten, 1996). Another treats it as the representation of intentional states, that is, states which are directed on things in the world (Hutto, Herschbach, & Southgate, 2011; Povinelli, Eddy, Hobson, & Tomasello, 1996). Representing propositional attitudes qualifies as mindreading on both of these accounts, since propositional attitudes are intentional states, and causally efficacious unobservables. But neither account requires mindreaders to represent propositional attitudes as such: one might represent a causally efficacious unobservable or intentional state without representing it as a propositional attitude.

In this connection, it's been suggested that there may be various forms of 'minimal' mindreading. These alternative forms of mindreading are *minimal* in the sense that they do not involve 'metarepresentation'. That is, they do not involve representing representational states as such. A representational state is a state with semantic properties – notably, having a content apt for accuracy and inaccuracy. The content of a state is what it represents. For instance, my belief that my house is white represents *that my house is white*. My belief is accurate just in case my house is white. To represent a representational state as such is to represent it as having these semantic properties. We do this when we ascribe propositional attitudes; minimal mindreaders, by definition, do not.

This distinction between minimal and metarepresentational mindreading has been explored in a number of places. One prominent account of minimal mindreading, for instance, suggests that it involves exploiting a 'minimal theory of mind' which defines non-representational surrogates for our mental state concepts. These are concepts of relations agents stand in to objects in their environment. The simplicity gained by representing mental states as relations rather than as representations comes at a cost. There are 'signature limits'

⁶ I use the term 'propositional attitude' to pick out mental states having these properties; I don't intend by this use that these are the only states expressing propositions. Whether other states express propositions turns in part on what propositions are, a question which is orthogonal to my concerns. If, for instance, propositions are sets of possible worlds, then presumably all representational states will express propositions. In what follows, I use the more neutral term 'content' to refer to what is expressed by a representational state.

on the abilities of minimal mindreaders. That is to say, there are certain mental states they cannot ascribe, but which ‘full-blown’ mindreaders can (Apperly & Butterfill, 2009; Butterfill & Apperly, 2013).⁷ In a similar vein, José-Luis Bermúdez (2011) distinguishes ‘perceptual mindreading’ from propositional attitude mindreading, where the former, again, involves representing relations between agents and objects or states of affairs, and the latter involves the ascription of propositional attitudes. And Josep Call and Michael Tomasello (2008) have proposed that chimpanzees understand others using a ‘perception-goal psychology’, distinct from and simpler than the belief-desire psychology with which adult humans operate.

In each case, the suggestion is that minimal mindreaders differ from ‘full-blown’ mindreaders in construing mental states as relations between agents and things in the environment. This is contrasted in each case with ascribing representational states, that is, states with contents apt for being accurate or inaccurate. But more specifically, in each case, the contrast is drawn between minimal mindreading and the variety of mindreading engaged in by adult humans – propositional attitude mindreading. For instance, Butterfill and Apperly (2013) write that ‘full-blown’ mindreading is more challenging than minimal mindreading because propositional attitudes have a number of complex properties. For instance, they interact with one another in uncodifiably complex ways, and have arbitrarily nestable contents – such that one can think things like ‘Bill believes that Priya hopes the Bobcats will win the game.’ The suggestion is that full-blown mindreading involves appreciating this, and representing mental states as having these features. Similarly, Bermúdez (2003b, 2018) highlights that, unlike perceptual states, propositional attitudes do not impact behavior directly, but only in concert with one’s other beliefs and desires, in virtue of the logical, structural relations holding between these attitudes. As a result, he claims, a propositional attitude mindreader must be able to represent the logical relations between attitudes, and ‘work out how the relevant beliefs or other propositional attitudes will feed into action’ by ‘recreating’ the target’s reasoning.

These are no doubt features of the propositional attitudes, and correspondingly features of the propositional attitude mindreading often engaged in by adult humans. But this should not lead us to think that propositional attitude mindreading is the only alternative to minimal mindreading, or that representing representational states as such necessarily involves representing them as having these other properties. Just as a minimal mindreader fails to

⁷ Andrews (2018) also proposes that animal mindreading may make use of these surrogate concepts, but does not endorse all elements of Butterfill and Apperly’s view.

represent the semantic properties of mental states, a metarepresentational mindreader's grasp of the features of mental states might be partial. There is surely a region of logical space here that has so far been overlooked: a region occupied by creatures who know that others represent and misrepresent the world, but who do not take them to have propositional attitudes as such.

As a first step in exploring this region, consider that to ascribe representational states to others, we represent that they have representations which express an accuracy-evaluable content. Any representation which expresses something accuracy-evaluable expresses it in some representational format. By a representational format I mean something like a sentence, a diagram, a map or a picture. The same content can be expressed in different formats. For instance, 'my house is white' expresses that my house is white, but a picture of my house might also express this. What is expressed by the sentence 'average monthly rainfall in New York doubled in the year to April 2018' could also be expressed by a diagram. I take the differences here to be differences in representational format, rather than in content.

Representational states themselves can exploit a range of representational formats. This is significant, since representational formats differ in certain important ways, with the result that some are better suited than others to the expression of certain contents, and some have greater expressive power than others. As David Marr (1982) notes, the representational format one employs 'determines what information is made explicit and hence what is pushed further into the background, and it has a far reaching effect on the ease and difficulty with which operations may subsequently be carried out on that information'.⁸ The representational format exploited in a cognitive task may consequently have substantial effects on one's cognitive capacities and behavior. So, the idea that a thought has one representational format rather than another can be understood functionally (Camp, 2007, p. 174).

To take an example, Jacob Beck (2015) has argued that many animals use 'analogue magnitude representations' to represent quantities. These are primitive systems for representing quantity that do not rely on grasping units of measurement or number systems. A system of dashes in which --- stands for 3 and ----- for 5 is an example. An effect of representing quantities using analogue magnitude representations is that one's ability to discriminate between two quantities diminishes as the ratio between those quantities approaches 1:1. Discriminating ----- from ----- is harder than discriminating --- from

⁸ (Reprinted in Cummins & Cummins, 2000, p. 70).

----. So, this is a case where the format of thought has a demonstrable effect on one's cognitive capacities and resulting behavior.

In the light of this, we can ask: when a mindreader ascribes a representational state, what representational format do they take it to have? There are a number of reasons to think that in ascribing propositional attitudes, we take them to have a linguistic format. Most obviously, when we ascribe propositional attitudes, we use sentences to express their contents. More importantly, we take the attitudes to stand in logical relations which are effectively mirrored by the sentences we use to express them, indicating that we take them to have the same formal structure. For this reason, Bermúdez (2003b, 2011, 2018) has argued that the contents of propositional attitudes *must* be represented linguistically: only a linguistic format can make perspicuous the logical relations holding between distinct thoughts in virtue of which they give rise to the behavior they do. Similarly, Butterfill and Apperly note (2013), propositional attitudes have nestable contents, meaning that a person think things like 'Sally hopes that Bob knows that Pesh wants to go to university.' But these nested contents are of a kind which plausibly can be captured only in a linguistic format (Camp, 2007). So, in taking this kind of thought to be possible, we commit to thought having a linguistic format. More generally, we tend to assume that thought is topic neutral – such that one can think virtually anything about anything. For this to be the case, though, we must take the format of thought itself to be topic neutral – that is, to be sufficiently expressively powerful to capture the range of possible thoughts. Language is the only plausible candidate for having this expressive power; as will become clear in the next section, other formats are far more expressively limited.

Although we take propositional attitudes to be linguistic in format, there is no reason to think that this is necessary. One might treat representational states as having a non-linguistic representational format. Of course, this would mean failing to capture certain important features of those states – that they can have nested contents, or that they stand in all the inferential relations they do. Indeed, as should already be clear, ascribing states with a non-linguistic representational format is likely to have a significant impact on the mental state ascriptions one could make, and the ease or difficulty with which one could make them. As a result, returning to the question with which I began, from the fact that great apes seem to represent representational states as such, it would be far too quick to conclude that their mindreading is substantially the same as our own.

3. Mindmapping

To illustrate this, I consider in this section what it would be to ascribe states with a map-like format, rather than a linguistic one. To begin with, consider the comparison between a language and what I shall call a ‘basic map’.⁹ I use this term to refer to a representational format which exploits a lexicon of syntactically simple icons denoting objects and properties, and the combinatorial principle of spatial isomorphism. According to this principle, placing two icons in a spatial relation on a map has the effect of representing that the referents of those icons stand in an isomorphic spatial relation in reality (Camp, 2007, p. 158).

Both languages and basic maps are combinatorial representational formats, in which discrete representational units are combined according to combinatorial rules. The semantic content of the resulting complex representations is a systematic function of the semantic content of the units and the way in which they have been combined (Camp, 2007, p. 154). But maps and language make use of very different combinatorial rules. In language, the combinatorial principles are abstract, relying on no physical similarity between the representation and what is represented. For instance, in the sentence ‘Boris is wise’, the syntactic relation of functional application is used to mirror the metaphysical relation of instantiation (Camp, 2007, p. 157). By contrast, the principle of spatial isomorphism exploited by maps imbues the spatial properties of a representation with semantic significance.

The result of employing such different combinatorial principles is that these formats differ in characteristic ways. In some ways, basic maps are more efficient and user-friendly than languages. First, they are informationally dense. Reproducing the representational content

⁹ I use this term to distinguish the maps I am talking about from augmented map-like formats, to be discussed briefly later. A basic map, as I use the term, is less powerful than the augmented map-like formats discussed by Elisabeth Camp (2007), and from those Frank Jackson and David Braddon-Mitchell (2007) have in mind when they argue that human thought might be map-like. In this section, I follow Camp (2007; see also Rescorla, 2009) in taking the fundamental combinatorial principle of a map to be spatial isomorphism rather than, say, predication (as proposed by, e.g. Casati & Varzi, 1999; Kulvicki, 2015). However, the claims that follow about the expressive limitations on maps should be acceptable even to those who adopt a predicational semantics, on which maps are conjunctions of predications of properties to locations. It’s less obvious that such a semantics can capture the idea that maps are ‘user-friendly’ in the ways to be outlined; following [redacted, under review], I take that to count against a predicational account.

of a basic map in sentential form is a difficult task, requiring a large and unwieldy set of sentences. Second, they wear their implications on their sleeves. For instance, if a map of the United States shows that Seattle is north of Portland and west of Spokane, it will also explicitly represent that Portland is south-west of Spokane, that Spokane is east of Portland, and so on. This information will be *cognitively transparent*, because it is explicitly represented in the map. By contrast, the sentence ‘Seattle is north of Portland and west of Spokane’ *implies* all that information, but does not explicitly represent it; someone using this sentence must do some work to figure out the rest. Finally, an effect of this is that making any alterations to a basic map is simple: moving any one icon on the map automatically updates all the represented relations between the represented objects, with the map’s coherence being automatically maintained (Camp, 2007, pp. 160–162).

These advantages come at the cost of flexibility. What a representational format can represent is determined by its combinatorial principle: it can represent things as standing in the relation represented by that principle. The combinatorial principles in language stand for very general relations like instantiation, which can relate almost anything; as a result, language has almost unlimited expressive power. By contrast, the combinatorial principle of a basic map is ‘semantically robust’, which is to say, it heavily constrains which contents a map can express (Camp, 2009, pp. 120–121). Whenever two icons are combined according to this principle, what is expressed is always something about the spatial relation between their referents.

Using only the resources of a basic map – that is, a lexicon of syntactically simple icons and the principle of spatial isomorphism – one cannot represent non-spatial relations between objects and properties such as *x is two years old*, *y is heavier than z*, or *o instantiates p*. Of course, if a suitably expansive lexicon is used, these resources can nevertheless go a long way. For instance, if a lexicon included distinct, syntactically simple icons for ‘green apple’ and ‘red apple’, a map could discriminate between there being a green apple at a location and there being a red apple at that location. But this would not amount to representing *that an object has a non-spatial property* – since the icons do not have formally separable elements corresponding to the object and the property. To put it another way, a map tokening syntactically simple red-apple and green-apple icons would not thereby represent that two things of the same kind were in the mapped region.

There are also limits on the quantificational information that can be represented in basic maps. For example, they cannot represent bare existential quantification. That is, they cannot say simply that something exists, without saying anything further about it. A map can

only represent that something exists by placing an icon *somewhere* on the map – and so, by representing that it exists in a particular location, and stands in certain spatial relations to other objects (Camp, 2007, p. 165). Representing identity on a map is another challenge. Whilst language permits us to say, ‘Bruce is Batman’, a map of Gotham City which had a Bruce icon and a Batman icon could not do the same. We might co-locate the Bruce and Batman icons, but the map itself provides no resources for representing that these two icons stand for one and the same person, rather than two co-located individuals.

The constraints I have been discussing affect what I’ve been calling ‘basic maps’. As Camp (2007) argues, maps can be augmented with notation and conventions which increase their expressive power, so these constraints do not apply absolutely to anything which we might be inclined to call a map. These more powerful, augmented formats effectively borrow representational resources from other formats, and are, in that sense, hybrid. One version of a hybrid map might employ syntactically complex symbols to represent non-spatial properties and relations. For instance, the colour of an icon might be used to represent the colour of its referent. The resulting icon would be syntactically complex, because it would have formally separable elements – colour and shape – each having semantic significance (Camp, 2007, p. 166, n. 30). Additional notation could also support the representation of object identities. For instance, we might say that where two co-located symbols are linked by a ‘=’ sign, they represent a single object. However, these additional resources make hybrid maps less cognitively transparent, and less easy to update. Just as efficiency in basic maps is bought in the coin of expressive power, expressive power in hybrid maps is bought in the coin of efficiency.

Now, having outlined these comparisons between a linguistic format and the basic map format, let us introduce the term ‘mindmapper’ to denote a creature who represents representational states as having a basic map format. Rather than taking a belief, say, to have a linguistic content, a mindmapper takes a belief to be a ‘map of the neighbouring space by which we steer’ (Ramsey, 1931).¹⁰ In other respects, let’s say, their picture of belief is similar to ours – that is, they take a belief to be a state which aims accurately to represent the world, and they expect others to act as though their beliefs are true. The difference is simply that they take beliefs to represent in a map-like rather than sentence-like way.

¹⁰ Of course, this appropriation of Ramsey’s phrase is undoubtedly more literal than he intended.

Since maps make a certain kind of information explicit at the cost of expressive power, mindmapping would give rise to a distinctive and limited set of mindreading abilities. In some ways, mindmapping would have its advantages. As Bermúdez (2003b, 2018) has argued, beliefs do not often give rise to behavior individually, and if one is usefully to ascribe more than one, one must be capable of recognizing the connections between them. By representing someone not as having a number of discrete sententially structured beliefs, but as having a single map-like representation of their environment, one might capture in a single ascription what a propositional attitude mindreader would treat as a number of distinct beliefs. Moreover, since maps make explicit the spatial relations holding between any objects they represent, this would take some of the cognitive work out of determining the relations between these beliefs.

As well as this, updating belief ascriptions would be somewhat simpler for a mindmapper. Suppose a mindmapper, Nila, represents a target, Robin, as having the map-like belief in Figure 1. The map represents Nila and Robin in their shared environment, separated by a stream. A food item, represented by the bubble icon, is represented as being on Robin's side of the river. Now suppose that Nila moves the food over to her side of the river, and she knows that Robin was watching as she did so. To update her representation of Robin's belief, Nila moves the food icon to the relevant location on the map (Figure 2). As well as updating the represented location of the food, this automatically updates the represented spatial relations between the food and everything else on the map and maintains the coherence of the map. By contrast, if Nila were a propositional attitude mindreader, ascribing beliefs with a linguistic format, things would be more complicated. The sentence which gave the food's location would need to be updated, and a further process would need to check this revised sentence for consistency with any other beliefs ascribed to Robin. This additional processing introduces a risk of error where none exists for a mindmapper.

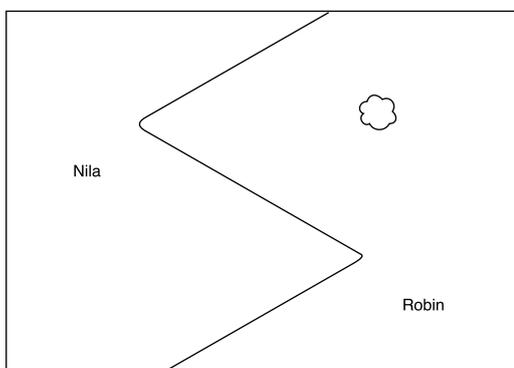


Figure 1

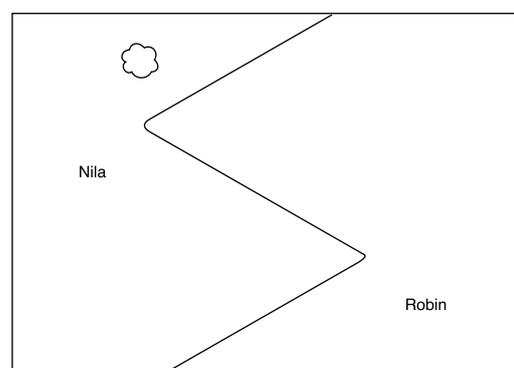


Figure 2

However, the range of mental state ascriptions a mindmapper can make is constrained by the expressive limitations on maps, outlined above. As a mindmapper, Nila should be able to ascribe mental states whose content concerns the spatial arrangement of objects and properties, but not those whose content concerns other properties and relations. So, she can ascribe a belief about the location of some food, say – but since she treats beliefs as having a basic map format, she cannot ascribe beliefs about other properties of the food, such as its colour or how long it's been there.

She will also be unable to ascribe beliefs involving bare existential quantification. So, for instance, suppose Nila moves the food to a new location, without Robin noticing. Nila takes Robin to have a belief whose content is expressed by the map in Figure 2 – which is now inaccurate. But now suppose that Nila observes Robin go to the place where she (Robin) believes the food to be located, and discover that there is no food there. Nila will now have to revise her belief ascription, since Robin no longer believes that there is food at that location. The most obvious result would be to ascribe a belief whose content is expressed by the map in Figure 3. There is no food icon on this map, because at this point, let's say, Robin may have no idea where the food is. Even if she did, Nila has no idea where Robin thinks it is, and so no reason to ascribe a belief to Robin about the food's location. Unless a food icon is placed on the map, the map no longer represents that any food exists.¹¹ Yet Robin may well believe that the food still exists. It's just that she has no idea where it is. Because maps cannot capture this sort of bare existential quantification, Nila's belief ascription in this situation will be inadequate.

¹¹ On many accounts, it represents that there is *no* food in the region. According to what Michael Rescorla (2009) calls 'the absence intuition' – according to which, if a map's lexicon has an icon which does not appear on the map, the map represents that the thing denoted by that icon is absent from the represented region. Rescorla and several others (see Casati & Varzi, 1999; Kulvicki, 2015) take this to be a semantic feature of maps. Elliot Sober (1976) and Elizabeth Camp (2007), on the other hand, take it to be merely pragmatic.

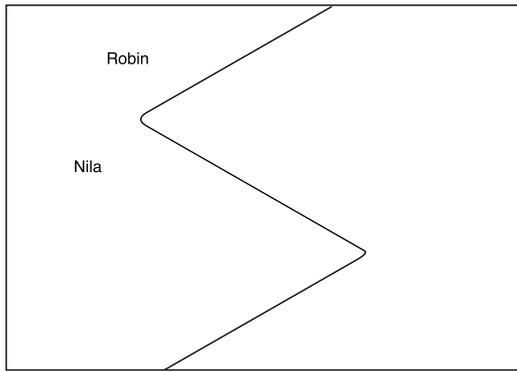


Figure 3

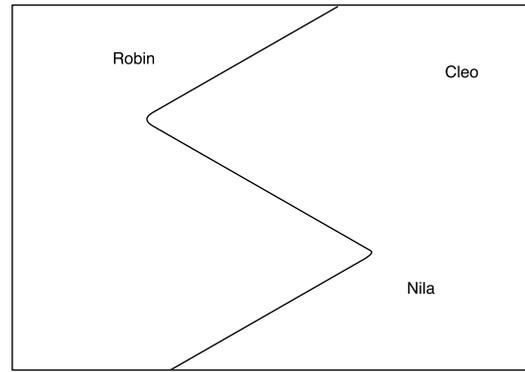


Figure 4

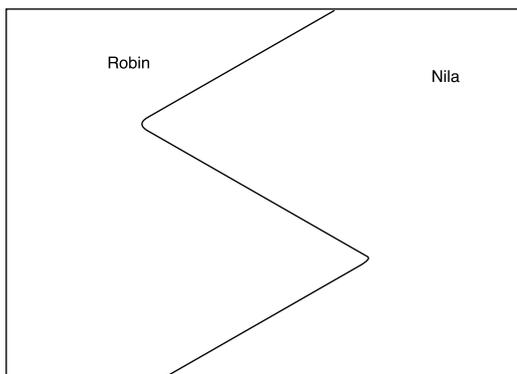


Figure 5

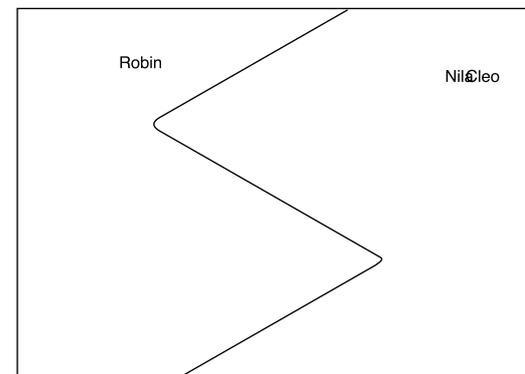


Figure 6

Finally, Nila will be unable to represent beliefs with content of the form $x=y$. So, suppose Nila has an alter-ego, Cleo, and Robin does not initially know that they are one and the same. Robin might come to believe, falsely, that Nila is in one place and Cleo in another. Nila might ascribe to Robin a belief with the content expressed by Figure 4. But now, suppose that Robin learns the truth: that Nila and Cleo are identical. Nila will be unable to ascribe this belief to Robin, since a basic map cannot express that two icons co-refer. Nila's only options would be to ascribe a belief with the content expressed in Figure 5, from which one of the icons is removed, or Figure 6, on which the icons are co-located. The first has the advantage of correctly capturing Robin's belief about the overall number of objects – but does not at all capture what it is she has learned. The second comes closer to capturing what was learned, but incorrectly ascribes to Robin a belief that there is one more object than there is.¹²

¹² Despite this, there are some beliefs about identity mindmappers can ascribe – those whose content is not of the form $x=y$. For instance, in the example above, using two icons for a single object enabled Nila to ascribe to Robin a belief that there were two distinct things. Similarly, in a case where a target mistakes the F before her for a G, a map with a G icon in the relevant location would enable one to capture the target's belief that there is a G there. This is one area in which mindmappers' abilities outstrip those of Butterfill and Apperly's 'minimal

In this section, I've been describing what it would be to represent others' beliefs as having a map-like format, rather than a linguistic one. I've argued that this would provide an effective means of capturing others' beliefs about the spatial distribution of objects, but otherwise would be severely limited. My purpose in making this argument has been to illustrate that the format one takes representational states to have substantially affects the range of mental state attributions one can make, and the ease or difficulty with which one can make them. So, for any representational format – including the augmented map formats I mentioned briefly above – we might describe a distinctive variety of mindreading which results from taking representational states to have that format. There is consequently, within the category of metarepresentational mindreading, significant scope for variation. So, from the fact that a creature represents representational states as such, we cannot conclude that its mindreading abilities do not differ substantially from our own.

4. Are Great Apes Mindmappers?

So far, I have been defending the general claim that there are varieties of metarepresentational mindreading falling short of the 'full-blown' propositional attitude mindreading of adult humans. I've argued that more limited forms of mindreading result from taking representational states to have a non-linguistic representational format, since the expressive limitations of such formats impose constraints on the mindreading tasks one can perform. In this final section, I consider the implications of this argument for the particular case of great ape mindreading. I argue that, consistently with the evidence, great apes may be mindmappers, but that this is a hypothesis which further empirical work could rule out. I argue that this is a project worth undertaking, and offer some suggestions about how this might be done.

The best evidence that great apes represent representational states as such comes from a pair of false belief tests in which they are faced with an agent who has a false belief about the location of an object (Buttelmann, Buttelmann, Carpenter, Call, & Tomasello, 2017; Krupenye et al., 2016). Since in both cases the belief ascribed concerns the location of an object, it is a belief which could be ascribed by either a mindmapper or a 'linguistic mindreader' – that is, one who takes mental states to have a linguistic format. So, the results of these studies are consistent with the hypothesis that great apes are mindmappers. Similarly, much of the

mindreaders', who are unable to represent such mistakes. (Butterfill & Apperly, 2013; see also Carruthers, 2015)

evidence that great apes and other animals ascribe knowledge is drawn from tasks involving the ascription of knowledge about the location of a food item (Bugnyar, 2011; Clayton et al., 2007; Hare et al., 2001; Martcorena et al., 2011). This kind of test cannot discriminate mindmappers from linguistic mindreaders, since both should be able to ascribe states with this kind of content. So, these results are consistent with an explanation in terms of either linguistic mindreading or mindmapping.

Naturally, my point is not that great apes *are* mindmappers. Rather, it is that this is an area in which our evidence underdetermines the theoretical possibilities: we do not know what format great apes take mental states to have. That great apes are mindmappers is therefore an epistemically open possibility, one which further empirical work might rule out. Of course, one might reasonably ask at this point whether this is a possibility which is worth ruling out. Not every possibility that is consistent with a body of evidence ought to be taken seriously. I propose, though, that there are a number of reasons to take this possibility seriously.

First, there is some limited empirical support for the idea that great apes are mindmappers. Current evidence indicates that great apes are incapable of level-two visual perspective taking (Karg, Schmelz, Call, & Tomasello, 2016). Level-two visual perspective taking is the ability to represent not merely *which* objects another individual can see,¹³ but how those objects *look* to her – and to appreciate that the very same objects may look different from her perspective than they do from one’s own (Flavell, Everett, Croft, & Flavell, 1981; Salatas & Flavell, 1976). Level-two visual perspective taking is often taken to be closely connected to the false belief test – both rely on appreciating that another person can have ‘a mistaken perspective’ (Karg et al., 2016). So, it might be surprising that great apes have passed false belief tests, whilst failing level-two visual perspective taking tasks.

The mindmapping hypothesis can provide an explanation of this pattern. As I mentioned briefly in §2, the idea of mindmapping can be applied to content-bearing states other than belief – so, a mindmapper might represent the content of visual perception as map-like. A creature who did this would be capable of level-one, but not level-two visual perspective taking – because a basic map can represent where things are, but not what they look like. To see this, consider a situation in which a mindmapper M and her target T sit on opposite sides of a cube. The face of the cube pointing toward M is green; the face pointing towards T is red. M can represent that T sees the cube, by ascribing to her a perceptual state whose content is

¹³ Representing *which* things others can see, and recognising that these things may differ from the things one can see oneself, is level-one visual perspective taking.

expressed by a map with the ‘cube’ icon at the relevant location. But she can’t represent that T sees this cube *as red*. This would require something syntactically more complex – an icon denoting the cube, which could be placed at the relevant location to represent where the cube was, and which could be modified in some other way to represent its being red. A basic map lacks these resources. Note that, even if the map’s lexicon contained a syntactically simple ‘red-cube’ icon, placing it on the map would not amount to level-two visual perspective taking. If the ‘red-cube’ icon denotes the cube, this again simply represents that T can see the object – but does not speak to how she sees it. If on the other hand, the ‘red-cube’ icon instead denotes the property ‘red-cubeness’, the map will represent that T sees red-cubeness here. But this leaves it an open question whether T sees this as a property of the very same object M is looking at. Level-two visual perspective taking, though, involves recognising that I and another see *the very same object* in different ways.

Despite this inability to engage in level-two visual perspective taking, though, a mindmapper can nevertheless ascribe some false beliefs – those whose contents can be ascribed in a map like format, which includes those beliefs ascribed in standard false belief tests. So, the hypothesis that animals are mindmappers explains this aspect of the evidence concerning great ape mindreading.

Second, there is already good reason to think that non-human cognition exploits a variety of representational formats. For instance, it’s been suggested that non-human animals might represent dominance hierarchies using a diagrammatic format (Camp, 2009); that they navigate using map-like representations (see Rescorla, 2018); that imagistic representations may underpin their physical reasoning (Gauker, 2018); and that they represent quantities using analogue magnitude representations (Beck, 2015). Michael Tomasello (2014) has proposed that great ape cognition exclusively employs a system of iconic representation which ‘prefigures’ the linguistic representation exploited by human thought. Against this theoretical landscape, it does not seem far-fetched to suppose that great apes might take others’ mental states to have a map-like format, rather than a linguistic one. And on the flipside, there is no *prima facie* reason to suppose that they should take representational states to have a linguistic format.

Finally, the alternatives to taking this possibility seriously are unattractive. We might simply suspend belief about what format great apes take representational states to have. But this would impair our ability to predict which mental states they could ascribe, and would leave us with a picture of great ape mindreading which was in an important respect incomplete. Alternatively, we might conclude that great apes take representational states to have a linguistic format, as we do.

But given the connection outlined above between taking mental states to have a linguistic format and representing propositional attitudes as such, this would involve ascribing to great apes an understanding of the mind considerably more sophisticated than the evidence mandates. This would be in violation of Morgan's Canon – the methodological principle admonishing us not to explain animal behavior in terms of a more complex psychological capacity when a simpler one will do (Morgan, 1894). Of course, that principle is rightly controversial (Fitzpatrick, 2008; Sober, 2005). But this move would also violate Evidentialism – the principle that we should not endorse an explanation of an animal's behavior in terms of a cognitive process X if our evidence gives us no reason to prefer it to an explanation in terms of another process Y (Fitzpatrick, 2008).

Whether great apes are mindmappers is a question we might begin to answer by implementing mindreading tasks which depart from the simple object-location paradigm mentioned above. States whose contents concern objects' non-spatial features are the most obvious to consider. But to show that a subject ascribes a belief about non-spatial properties is more difficult to establish than it might initially appear. Consider the following study conducted with children. Subjects observe as an experimenter demonstrates for the benefit of a target that an object makes a rattling noise. The experimenter asks the target 'can you do that?' and offers her a choice between two further objects. Children expect that the target will choose the object that looks similar to the first, rattling object, even though the children know, having been familiarised with the objects, that it is the other (dissimilar) object which rattles (Scott, Baillargeon, Song, & Leslie, 2010). This is taken to show that the children can ascribe beliefs about non-spatial properties. But mindmappers could pass this test by placing a 'rattle' icon on the map modelling the target's belief, thereby ascribing the belief 'there is a rattle at that location', rather than 'that object is a rattle'. Using a suitably expansive lexicon of syntactically simple icons, a mindmapper could make many of the same behavioural predictions as a linguistic mindreader who ascribes beliefs about the non-spatial properties of objects.

A task with a temporal dimension, in which an object's properties change over time, might be more effective. Here is an illustrative example. In a study on episodic memory in great apes, subjects observed two food items being hidden. One of the items was a grape; the other a piece of frozen juice, which they preferred to the grape, but which would melt after a certain amount of time. Then, after either five minutes or an hour, they were given the opportunity to choose one of the hiding spots. They chose the frozen juice significantly more after five minutes, and the grape significantly more after an hour (Martin-Ordas, Haun,

Colmenares, & Call, 2010). This indicates, among other things, that they know that over time, the frozen juice melts and becomes unobtainable. Now we might approach the question whether apes are mindmappers by considering whether they attribute knowledge of that kind to others – knowledge that frozen juice melts and becomes unobtainable. One way to explore this would be to investigate how well apes predict the behaviour of other individuals who are either ignorant or knowledgeable about the behaviour of frozen ice. A linguistic mindreader should be capable of representing the difference between these individuals: one knows, and the other does not, that frozen ice melts. A mindmapper should not, since a basic map cannot represent that something has non-spatial properties which change over time.

We might also investigate whether apes can ascribe beliefs about objects whose location the target individual does not represent – things, for instance, like ‘the ball is blue’. Suppose a subject watches as an agent A plays with a ball whose colour changes from blue to yellow and back each time it is tossed into the air. From time to time, A puts the ball in a box and puts a lid on it. While the box is closed, a bell sounds. At this moment, she presses either a yellow or a blue button corresponding to the colour she thinks the ball is. In the test trial, things proceed in the same way, but after putting the ball in the box, A leaves the room. Whilst she is gone, an experimenter takes the ball from the box, changes its colour, and replaces it. She now either takes the box away and then returns with it (box-present condition) or she takes the box away and returns without it (box-absent condition). A returns, and the bell sounds. Anticipatory looking should reveal whether subjects expect her to press the yellow or blue button. The critical question is whether the presence or absence of the box makes a difference. Assuming the task demands are not overwhelming, a linguistic mindreader ought to anticipate that A will press the button *not* corresponding to the present colour of the ball. But a mindmapper ought to anticipate this only in the box-present condition. In the box-absent condition, she should be indifferent – since A does not know or represent the location of the ball, and mindmappers cannot ascribe beliefs about objects without ascribing beliefs about their location.

Of course, there are limits to the type of investigation I’m proposing. Tests of this kind might conclusively establish the negative result that a subject was not a mindmapper. But no such test could conclusively establish the positive result that a subject *was* a mindmapper. Anything within the capacities of a mindmapper should also be possible for a mindreader using a more powerful representational format. As such, there is unlikely to be a behavioural test on which a positive performance is conclusive evidence of mindmapping. Conversely, there will be no single test in which failure is conclusive evidence of mindmapping, since subjects might

fail a test for many reasons, including excessive task demands. Consequently, it may be that no single test could definitively diagnose a subject as a mindmapper. Nevertheless, the question whether great apes are mindmappers is an empirically tractable one. It is a question which cannot be answered using a single test – but what matters is the possibility of a battery of tests revealing a pattern of behaviour best explained by the mindmapping hypothesis.

One might object that even such a rich pattern of evidence might be multiply interpretable, because just as maps can be extended in various ways, languages can be arbitrarily restricted. By restricting the available predicates in a language's lexicon, we might arrive at a representational format with precisely the same expressive power as a basic map, but which would nevertheless be a language, in virtue of employing linguistic combinatorial principles. Thus, even if we discovered a creature whose mindreading was limited to the ascription of mental states with contents carrying concrete information about spatial arrangements, this evidence would support two hypotheses equally well. Either the creature uses a basic map format to model mental states, or they use an arbitrarily restricted language. Consequently, any pattern of behaviour indicative of mindmapping is also consistent with some version of linguistic mindreading.¹⁴

In reply, we might dispute whether such an arbitrarily restricted form of linguistic mindreading would be empirically equivalent to mindmapping – since, as already noted, maps and language differ in respects other than expressive power. But even if this claim is granted, it should not be concerning, since fitting the data is not the only theoretical virtue. There's a natural understanding of the nature of cartographic representation from which the distinctive pattern of efficiency and limitation discussed in this paper unfolds. In the context of mindmapping, it readily and independently gives rise to empirical predictions. But there is no natural understanding of a language exhibiting this same pattern. Any empirically equivalent linguistic mindreading hypothesis must be produced by constructing a language in an ad hoc fashion to fit the pattern maps produce. It will give rise to certain predictions only because it has been designed to give rise to *precisely those* predictions, and has no independent predictive power. Being ad hoc and lacking in predictive power are both theoretical vices. So, even if there were a situation in which the evidence did not conclusively tell between a mindmapping hypothesis and a restricted linguistic mindreading hypothesis, it might nevertheless be reasonable to favour one over the other.

¹⁴ Acknowledgment removed for anonymous review.

5. Conclusion

In this paper, I have been arguing that mindreaders might vary with respect to the format they take representational states to have. This is significant, since the format one takes mental states to have imposes substantial constraints on one's mindreading abilities. I have illustrated this by considering the possibility of 'mindmapping' – that is, taking mental states to have the format of a basic map. I showed that this would result in a form of mindreading which was constrained by the expressive power of a basic map. In brief, it would provide an effective way to represent mental states whose contents concern the spatial arrangement of objects and properties, but would not support the ascription of mental states concerning non-spatial properties, bare existential quantification, or identity. The same idea might be applied more generally: for any format, we can describe a characteristic form of mindreading which is the result of taking mental states to have that format. So, the argument here reveals a vast space for theorising about forms of mindreading other than our own.

This has concrete significance for the emerging picture of great ape mindreading – since it demonstrates that, from the fact that great apes pass false belief tests, we cannot conclude that great ape mindreading and human mindreading are substantially the same. Consistently with what we know, great apes could be mindmappers. If they are, then their mindreading abilities will be considerably less powerful than our own, in a systematic and predictable way. I have not, however, argued that great apes *are* mindmappers; rather, I have been arguing that the evidence underdetermines what format they take mental states to have. Further empirical work could support taking a stand one way or the other on this question. This work, I have argued, is worth undertaking. Given that one's mindreading abilities are in part determined by the format one takes mental states to have, any story about animal mindreading which is silent on that question will be incomplete.

References

- Andrews, K. (2012). *Do Apes Read Minds? Toward a New Folk Psychology*. Cambridge, Massachusetts: MIT Press.
- Andrews, K. (2018). Do chimpanzees reason about belief? In K. Andrews & J. Beck (Eds.), *The Routledge Handbook of Philosophy of Animal Minds* 2. Abingdon: Routledge.
- Apperly, I. A., & Butterfill, S. A. (2009). Do humans have two systems to track beliefs and belief-like states? *Psychological Review*, 116(4), 953–970.
- Beck, J. (2015). Analogue magnitude representations: A philosophical introduction. *British Journal*

- for the *Philosophy of Science*, 66(4), 829–855. <http://doi.org/10.1093/bjps/axu014>
- Bermúdez, J. L. (2003a). The domain of folk psychology. *Royal Institute of Philosophy Supplement*, 53(1), 25–48. <http://doi.org/10.1017/S1358246100008250>
- Bermúdez, J. L. (2003b). *Thinking Without Words*. Oxford: Oxford University Press.
- Bermúdez, J. L. (2011). The force-field puzzle and mindreading in non-human primates. *Review of Philosophy and Psychology*, 2(3), 397–410. <http://doi.org/10.1007/s13164-011-0077-9>
- Bermúdez, J. L. (2018). Can nonlinguistic animals think about thinking? In K. Andrews & J. Beck (Eds.), *The Routledge Handbook of Philosophy of Animal Minds*. Abingdon: Routledge.
- Braddon-Mitchell, D., & Jackson, F. (2007). *The Philosophy of Mind and Cognition*. London: Blackwell.
- Bugnyar, T. (2011). Knower-guesser differentiation in ravens: others' viewpoints matter. *Proceedings of the Royal Society B: Biological Sciences*, 278(1705), 634–640. <http://doi.org/10.1098/rspb.2010.1514>
- Bugnyar, T., Reber, S. A., & Buckner, C. (2016). Ravens attribute visual access to unseen competitors. *Nature Communications*, 7, 10506. <http://doi.org/10.1038/ncomms10506>
- Buttelmann, D., Buttelmann, F., Carpenter, M., Call, J., & Tomasello, M. (2017). Great apes distinguish true from false beliefs in an interactive helping task. *PLoS ONE*, 12(4). <http://doi.org/10.1371/journal.pone.0173793>
- Butterfill, S. A., & Apperly, I. A. (2013). How to construct a minimal theory of mind. *Mind & Language*, 28(5), 606–637.
- Call, J., & Tomasello, M. (2008). Does the chimpanzee have a theory of mind? 30 years later. *Trends in Cognitive Sciences*, 12(5), 187–192.
- Camp, E. (2007). Thinking with maps. *Philosophical Perspectives*, 21, 145–182.
- Camp, E. (2009). A language of baboon thought? In R. W. Lurz (Ed.), *The Philosophy of Animal Minds* (pp. 108–127). Cambridge: Cambridge University Press.
- Carruthers, P. (2015). Two systems for mindreading? *Review of Philosophy and Psychology*, 6(1). <http://doi.org/10.1007/s13164-015-0259-y>
- Casati, R., & Varzi, A. (1999). *Parts and Places*. Cambridge, Massachusetts: MIT Press.
- Clayton, N. S., Dally, J. M., & Emery, N. 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(February), 507–522. <http://doi.org/10.1098/rstb.2006.1992>
- Cummins, R., & Cummins, D. (2000). *Minds, Brains and Computers: The Foundations of Cognitive Science - An Anthology*. Oxford: Blackwell.
- Fitzpatrick, S. (2008). Doing away with morgan's canon. *Mind and Language*, 23(2), 224–246. <http://doi.org/10.1111/j.1468-0017.2007.00338.x>
- Flavell, J. H., Everett, B., Croft, K., & Flavell, E. (1981). Young children's knowledge about visual perception: Further evidence for the Level 1-Level 2 distinction. *Developmental Psychology*, 17(1), 99–103. <http://doi.org/10.1037/0012-1649.17.1.99>
- Gauker, C. (2018). Visual imagery in the thought of monkeys and apes. In K. Andrews & J. Beck (Eds.), *The Routledge Handbook of Philosophy of Animal Minds*. Abingdon: Routledge.
- Hare, B., Call, J., Agnetta, B., & Tomasello, M. (2000). Chimpanzees know what conspecifics do and do not see. *Animal Behaviour*, 59(4), 771–785.
- Hare, B., Call, J., & Tomasello, M. (2001). Do chimpanzees know what conspecifics know? *Animal Behaviour*, 61(1), 139–151. <http://doi.org/10.1006/anbe.2000.1518>
- Hutto, D. D., Herschbach, M., & Southgate, V. (2011). Editorial: social cognition: mindreading and alternatives. *Review of Philosophy and Psychology*, 2(3), 375–395. <http://doi.org/10.1007/s13164-011-0073-0>
- Karg, K., Schmelz, M., Call, J., & Tomasello, M. (2016). Differing views: Can chimpanzees do Level 2 perspective-taking? *Animal Cognition*, 19(3). <http://doi.org/10.1007/s10071-016-0956-7>
- Krachun, C., Carpenter, M., Call, J., & Tomasello, M. (2009). A competitive nonverbal false

- belief task for children and apes. *Developmental Science*, 12(4), 521–535.
<http://doi.org/10.1111/j.1467-7687.2008.00793.x>
- Krupenye, C., Kano, F., Hirata, S., Call, J., & Tomasello, M. (2016). Great apes anticipate that other individuals will act according to false beliefs. *Science*, 354(6308), 110–114.
<http://doi.org/10.1126/science.aaf8110>
- Kulvicki, J. (2015). Maps, Pictures, and Predication. *Ergo*.
- Marr, D. (1982). *Vision: A Computational Investigation into the Human Representation and Processing of Visual Information*. San Francisco, CA: W. H. Freeman.
- Martcorena, D. C. W., Ruiz, A. M., Mukerji, C., Goddu, A., & Santos, L. R. (2011). Monkeys represent others' knowledge but not their beliefs. *Developmental Science*, 14(6), 1406–1416.
<http://doi.org/10.1111/j.1467-7687.2011.01085.x>
- Martin-Ordas, G., Haun, D., Colmenares, F., & Call, J. (2010). Keeping track of time: Evidence for episodic-like memory in great apes. *Animal Cognition*, 13(2), 331–340.
<http://doi.org/10.1007/s10071-009-0282-4>
- Morgan, C. L. (1894). *An Introduction to Comparative Psychology*. London: Walter Scott.
- Penn, D. C., & Povinelli, D. 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://doi.org/10.1098/rstb.2006.2023>
- Povinelli, D. J., Eddy, T. J., Hobson, R. P., & Tomasello, M. (1996). What young chimpanzees know about seeing. *Monographs of the Society for Research in Child Development*, 61(3), i+iii+v-vi+1-189.
- Ramsey, F. P. (1931). *The Foundations of Mathematics*. London: Kegan Paul.
- Rescorla, M. (2009). Predication and cartographic representation. *Synthese*, 169(1), 175–200.
- Rescorla, M. (2018). Maps in the head? In K. Andrews & J. Beck (Eds.), *The Routledge Handbook of Philosophy of Animal Minds*. Abingdon: Routledge.
- Salatas, H., & Flavell, J. H. (1976). Perspective taking: the development of two components of knowledge. *Child Development*, 47(1), 103–109.
- Scott, R. M., Baillargeon, R., Song, H., & Leslie, A. M. (2010). Attributing false beliefs about non-obvious properties at 18 months. *Cognitive Psychology*, 61(4), 366–395.
<http://doi.org/10.1016/j.cogpsych.2010.09.001>
- Sober, E. (1976). Mental representations. *Synthese*, 33(1), 101–148.
- Sober, E. (2005). Comparative psychology meets evolutionary biology: Morgan's canon and cladistic parsimony. In G. Mitman & L. Datson (Eds.), *Thinking With Animals: New Perspectives on Anthropomorphism* (pp. 85–99). New York: Columbia University Press.
- Tomasello, M. (2014). *A natural history of human thinking*. Harvard University Press.
<http://doi.org/10.1007/s13398-014-0173-7.2>
- Whiten, A. (1996). When does smart behaviour-reading become mind-reading? In P. Carruthers & P. K. Smith (Eds.), *Theories of Theories of Mind* (pp. 277–292). Cambridge: Cambridge University Press.