PENULTIMATE DRAFT

To appear in De Caro, Ferretti Marraffa (Eds) *Cartographies of the mind*, Springer 2007

THE BELIEFS OF MUTE ANIMALS

Simone Gozzano

1.INTRODUCTION

The issue of animal minds continues to engender philosophical debate.[1] Famously, Descartes considered animals without speech as *automata*, complex and fascinating machines with no internal lives.[2] Such a view shaped the subsequent intellectual *milieu* even after Darwin's theory, that closed the gap between humans and other species. For a long time it has been thought that the only hope for unveiling the capacity for thought of these creatures, thus bringing mental light into the dark matter, was by teaching them a form of language, and this position gave rise to a number of research projects. From the Kellogs in the 1930s to the Hayeses in the 1950s up to the Gardners (1970s) and the Premacks (1980-90s) much effort has been spent in teaching animals a symbolic system so as to check their cognitive abilities (Wallman 1992). Notwithstanding recent interesting findings (Savage-Rumbaugh 1986), the paradigmatic presupposition on which these attempts are founded has recently been challenged: does thought *necessarily* require a language? That is to say: is it possible to conceive non-linguistic behavior as driven and planned by thought-like states?

Considerations in favor of animal mentality have been advanced by many philosophers[3] and ethologists[4] but the general attitude remains one of skepticism even among ethologists,[5] the chief objection being that animals behavior could be explained by reference to complex series of stimula and responses. There are intermediate positions[6] that grant some form of proto-thought to animals other than humans, but some of those who favor the attribution of concepts to animals continue to assert that concepts are dependent on the capacity for symbol processing.[7]

It is revealing that philosophers, empirical scientists, psychologists and ethologists converge on the thought-language issue as this opens up the possibility of exchanging ideas and methods between the two groups of scholars. In the spirit of stressing this interaction, in this chapter I will firstly discuss the arguments against animal beliefs; I will then outline some general conditions for the individuation of functional states with content and for the attribution of intentional states to animals,

conditions stemming from the arguments against animals' beliefs; and finally, I will test these conditions against empirical cases.

2.THE ARGUMENTS AGAINST ANIMAL BELIEF

In a paper written many years ago Malcolm imagines a dog chasing a cat. The cat runs toward an oak tree but, at the last moment, swerves and disappears up a nearby maple. "The dog does not see this maneuver and on arriving at the oak tree [...] barks excitedly into the branch above" (1972, 13). According to Malcolm we can say that the dog *thinks* the cat is in the tree but we cannot say that he *has the thought* "the cat is in the tree." This idea has been attacked by Davidson who rejects the possibility of having thoughts without mastering a language. According to Davidson, there are three key preconditions for the proper attribution of intentional states: *i*) beliefs, and in general thoughts, attributed to anybody must reflect semantic opacity,[8] because such a logical feature has been considered to distinguish talk about propositional attitude from talk of other things (cf. 1985, 475); *ii)* semantic opacity requires the mastering of concepts, and *iii)* the attribution of concepts, in turn, presupposes a form of conceptual holism. Because it is neither possible to attribute concepts nor a holistic net of concepts to animals, it is not possible to ascribe semantically opaque thoughts to them; hence speechless animals cannot have beliefs and desires. Finally, Davidson argued that in order to have a belief one has to have the concept of a belief, that is, of a state that could be true or false. Now, even if some authors have disputed the importance of semantic opacity as a requirement for intentionality (e.g., Searle 1983), I share Davidson's idea that this feature is central in the analysis of intentional states.[9] Following Davidson, then, we can transform the question regarding the supposed thoughts of mute animals into whether we can make sense of opacity in contexts that exclude linguistic expressions.

Throughout his analysis of animal beliefs, Davidson is substantially guided by something like "Russell's Principle", according to which, in order to think about an object or to make a judgment about an object, one must know which object is in question (Russell 1912, 58; cf. Evans 1982, 65, 90 *et pass.*).[10] In Davidson, the kind of knowledge Russell's Principle is calling for amounts exclusively to *conceptual* knowledge. However, construed in this form, Russell's Principle seems too strong. Imagine observing two patches of red texture, one of which is darker than the other. Suppose you come to believe that the patch on your left is darker than the other. Yet, given the possibility that you lack the relevant concepts for the two shades of color, it would follow that you cannot be credited with the mentioned belief. Moreover, adopting Russell's Principle seems incorrect for a deeper reason. Belief is the basis of knowledge in that the latter can be taken, at least, as justified true belief. Nothing can be known if it is not believed in one way or another. Now, taking something like Russell's Principle as a precondition for beliefs' attribution begs the question of this notion of believing because the precondition is far more complex than the condition to be fixed.

To save this Principle, another route is viable: weakening its satisfaction conditions. In the case of the two patches of color, Russell's Principle could be

satisfied by *nonconceptual* discriminatory abilities (cf. Peacocke 1992). This weaker construal is the one given by Strawson (1971), according to which the knowledge of an object amounts to the ability to discriminate such an object when perceived at a time; to recognize it if presented; and to discriminate facts about it. Hence, we may say that the discriminatory abilities can be, at least, "nonconceptual"—as to perceptual discrimination and recognition—and "conceptual"—as in the case of knowledge of facts.[11] Many have forcefully argued against the notion of nonconceptual content.[12] In its place, these authors have proposed the notion of *demonstrative concepts*, one that plays the same role as that attributed to nonconceptual contents, without being holistic. Imagine me uttering: "This one is edible" in front of a piece of yellowish stuff. The demonstrative is referentially clear enough to bring about the belief in my interlocutor that something is edible here without her knowing what *kind* of thing it is that is edible. So the demonstrative concept does not allow forming any inferential connection with the particular in question, while it allows just partial connections to the concept of *edibility*. The example, however, might be dramatized: if every object we eat cannot be recognized (imagine a world where food changes continuously and we have to check it every time) then there would be no stable inferential liaisons we can use to fix the concept of *edibility*. On the other hand, the belief *this is edible* would allow for opaque substitution in the attribution practice. So, while demonstrative concepts may not meet the individuation conditions for concepts such as Davidson has in mind — which is how he arrived at the position that there is no clear answer to the question of whether the Ancients believed that the Earth was flat (cf. Davidson 1975, 168)— they are sufficient for crediting the individual with beliefs.

As I said, I think that Davidson is right in claiming that one of the standards for attributing thought must be opacity, so we should presume that having thoughts entails having referentially opaque thoughts based on elements that themselves need neither to be of a fully conceptual nature nor holistically individuated or connected. The general idea, then, is that a system can be credited with a perceptually fixed belief that *p* if, during the fixation of the belief, the elements included in *p* are either *i*, perceptually discriminated or recognized by the system, or *ii*, the system discriminates facts about them. For instance, if the belief that *p* is, in this particular case, the belief that *a* is *F*, then the belief that *p* can be attributed to the system if the system is able to discriminate the individual *a* and the property *F* according to one of the three possible ways mentioned. All the discriminatory abilities of a system form what I shall label the *epistemic window* of a system. We then may say that the beliefs that can be attributed to a system are those whose elements fall within the system's epistemic window.[13] Some systems have, as it were, a poor epistemic window: most animals cannot be credited with the ability to know facts about objects. In this sense I think that many discussions on animals' concepts are biased by the idea that concepts are all-or-nothing entities and that the standard of these entities are set by skilled adult human beings. But concepts should not be taken in this way: my concept of *atom* differs greatly from that of a physicist, and my four years' old son may get the notion of *constituent part* which is to my concept of *atom* as this is to the physicist's one.[14]

In order to ascribe perceptually fixed beliefs to non-speaking animals, the epistemic window is not sufficient. In fact, the epistemic window is not enough to warrant the sensitivity of the system to semantically opaque thoughts, and since semantic opacity helps in distinguishing talk about beliefs from other talks, it is important to show how the epistemic window and semantic opacity converge.[15]

In this respect, it is important to inquire what the philosophical significance of semantic opacity is. Even if knowing entails believing, both epistemic conditions share the semantic feature of not necessarily allowing co-referential substitution *salva veritate*. This stresses that truth is a much stronger notion than knowledge and belief: if you could get a truth, so to speak, by itself, then you would get all the possible expressions of that truth; whereas if you were able to isolate the pure knowledge or belief of a certain fact, you will not gain the whole truth about that fact, even if your knowledge is justified and your belief is true. This is because other possible true formulations of that fact are not necessarily included in the isolated knowing or belief. The distance that is present between knowledge and belief, on the one side, and truth on the other, marks our epistemic fallibility; that is, marks that we may be wrong about a matter of fact even if we are right to believe the very same fact ("It is not Tully that denounced Catiline, Cicero did!"). It is epistemic fallibility, the possibility of being mistaken and of being able to recognize such mistakes, that makes sense of intentional attribution.[16] It is not necessary, then, to take the problem of mistaking as an extra requirement for the proper attribution of intentional states, as Davidson believes: such a condition is already included in the semantic opacity requirement.

3.THE REQUIREMENTS FOR BELIEVING

According to the functional approach, systems have complex internal workings, framed in terms of states and processes functionally construed. In one of the deepest analyses of this approach, Loar (1980) considers functional states as a net of horizontal links. This net encompass perceptual inputs, internal workings—as inferences and the like—and behavioral outputs. Functional states are individuated through a process of theory construction, in which the internal states are analyzed, compared and contrasted with the overall behavior of the system. This horizontal net also has some vertical connections. Some functional states, in fact, tend to be reliably activated by non-mental conditions, so that we can say that these states are *associated* with certain truth-conditions even if they are *independent* from them.[17]

One of the most prominent tasks of the functional approach, then, is to interpret the internal states and processes that exist between an input, whether perceptual or conceptual, and an output, any kind of behavior, so as to individuate the *content* of these states and processes. Interpretation can be considered as that epistemic process that aims at individuating the contents of functional states. These contents, in turn, may have different functional status if they are embedded in epistemic states (e.g. beliefs) or conative ones (e.g. desires). These contents and their functional status determine the causal role played by the states themselves. However, when we wonder whether to attribute beliefs to animals, it is essential to

consider which are the features of this notion that we take to be determining. It seems to me that there are three different ways in which we can spell out the notion of *content*.[18]

3.1Contents of type 1

If a system discriminates a property, say *F*, then we assume that there is some internal process we can interpret as an *F*-detector. The resulting state can be considered a functional state whose content is "*F* is present".[19] This interpretation is based on the relation of *indication*, the simplest form of interpretation.[20] Consider a gauge that detects the presence of water in a tank. We can say that it *indicates* the presence of water in the tank. However, the gauge indicates what it does independently of the way in which its object is individuated: if it indicates water, then it indicates H_2O molecules in a liquid state or the kind of liquid I just had a glass of. An indicating relation then allows only for transparent readings.[21] As to the problem of predicting the system's behavior, it does not matter which reading one picks up: all the possible equivalent readings have the same predictive or explicative power. Strictly speaking, the detection system indicates that the mechanism itself is in a certain situation: the float in the tank is in a certain position if everything is properly working. Complete transparency follows. The consequence is that contents of type 1 fail to be associated with the appropriate truth conditions only in case the system is *malfunctioning*. For instance, a gauge might indicate the presence of water in the tank even in absence of water *only if* the gauge or its detection apparatus are malfunctioning. If everything is properly working, then the content indicates what it is associated with. So, type 1 contents are transparent and can be false only in case of malfunction. Hence, no intentional states can be attributed to systems endowed just with content of type 1.

3.2Contents of type 2

Type 2 contents, like those of type 1, are typically associated with some truth conditions. However, a malfunction of the system is not necessary in order to sever the connection between a given state in the system and some truth-conditions which are usually responsible for the system to be in that state. Consider Malcolm's dog again. When we say that the dog barks at the tree because, loosely speaking, the dog believes that the cat is in the tree, we are saying that the dog is in an epistemic state that we may interpret as having the content "the cat is in the tree". As in the case of content of type 1, the interpretation of the internal state of Malcolm's dog allows for other co-extensional readings too.[22] However, the correctness of the functional state's interpretation does not imply that the truth-conditions hold: it is possible that the cat is not in the tree even if nothing in the visual apparatus of the dog is malfunctioning. Let's be clear about this: you may have individuated the correct content attribution, in terms of the explanations and descriptions of the system's behavior and in terms of the reliable association that activates that state, without this content being *true* and without there being any malfunction in the system. In order

to correct a state with an erroneous type 1 content you should fix (at least a *part* of) the *structure* of the system; to correct a state with an erroneous type 2 content you should provide *more information* to the system itself. In case a system exhibits a behavior with respect to which the attribution of type 2 content is justified, we must postulate some *representational* function between stimulus and response, that is, some cognitive activity that mediates between sensory inputs and motor outputs. What this shows is that this kind of content does not allow for every possible substitution: at least some substitution may change the system's behavior – which is, as we saw when examining Davidson's standard for attributing mental content to an agent, exactly what we want: the opacity criteria is vindicated by the advocate of attributing mental activity to animals other than humans.

3.3Contents of type 3

Here we get to the notion of content as is customarily attributed to human beings. In this case, contents allow for opaque readings and their correct individuation establishes a weaker relation to truth conditions. If John believes that Cicero denounced Catiline this does not entail that John believes that Tully denounced Catiline because, even if Cicero *is* Tully, John may not have this information. Ascriptive talk is opaque in case of states with type 3 contents. This is so because, as is commonly known, intentional contents refer to their object through a "mode of presentation", and modes of presentation, namely the various ways through which it is possible to identify a given state of affair or event, are more fine-grained than their associated truth conditions. In case of type 3 content it is essential to pick up the correct mode of presentation in order to have a good prediction or explanation of the system's behavior. Finally, as in the case of states with type 2 content, the possibility of John's having an internal state whose content does not reflect a given state of affair does not necessarily entail any malfunctioning in John's epistemic or perceptual apparatus. Similarly, type 3 contents allow for opaque readings and can be false without the system having any malfunctioning.

How should we use these different notions of content? It is important to notice that attributing type 3 contents to a system entails that it would show different behaviors with respect to the elements that form the content of John's belief, if these were caught by different modes of presentation. John saying "Yes" to the question: "Did Cicero denounced Catiline?" does not entail John saying "Yes" to "Did Tully denounced Catiline?". Can something like this happen in the case of animals in which type 2 contents operate? The difference between type 3 contents and type 2 contents is a difference in the degree of granular fineness. The advocates of the thought-language identification, as Davidson and possibly Chomsky, have maintained that if a thought is not fine-grained as ours it is not a thought. They have excluded the possibility of thoughts of different grain. The possibility of type 2 contents shows that there is *at least* another option viable: contents of type 2 fix equivalence sets between ways of describing the elements that constitute the content of a given belief that are much larger than those fixed by contents of type 3 but not as comprehensive as contents of type 1. How can this difference be assessed? By the

capacity that systems have to adapting their behaviors to changes in the mode of presentation.

4.ATTRIBUTING CONTENTS TO MUTE ANIMALS

In order to find plausible candidates of intentional behavior in animals I will consider deceptive behaviors because these can be performed both through language and mute action. Among human beings there are a variety of kinds of deceptions, and even if deception seems an intentional notion in itself, there are some cases in which it is possible to deceive in a non-intentional way (cf. Chisholm and Feehan 1977). In order to reinforce the intentional interpretation of deception it should be conceived as voluntarily and goal-directed in that the deceptive behavior should not be an automatic response (such as camouflage or mimicry) and should not achieve its end by chance (the deceiving result must be brought about as driven by a goal). Let me clarify this by an example. In the Odyssey, Ulysses deceives his enemies by concealing his identity. He arrives in Ithaca dressed as a vagrant. His deception is voluntarily and goal-directed: it is voluntarily because it is not an automatic response; it is goal-directed because Ulysses' acting as a vagrant is aimed not at the pure exhibition of his deceiving ability, but as the result of a plan intended at the verification of his enemies' political and moral behavior. Does this case of deception establish a positive case for an intentional attribution? May we say that Ulysses' enemies *believe* that "a vagrant has arrived"? I think the answer should be positive. Ulysses' enemies can be credited with the belief "a vagrant has arrived" because i) the elements "vagrant", "Ulysses" and the property of *arriving* fall within their epistemic window and these elements can be described through functional states interpreted in terms of their various contents; ii) it would be possible to ascribe to Ulysses' enemies the content "Ulysses has arrived" that is contingently equivalent to the content of the supposed belief; iii) the attribution of a functional state interpreted as having one content does not imply the attribution of a functional state interpreted as having the other, therefore we may say that Ulysses' enemies *believe* that "a vagrant has arrived".

It should be stressed, thought, that invoking voluntariness and goal-directedness is not question begging on intentional content: these notions, as I have construed them here, can be attributed by using type 1 contents while intentional notions require at least contents of type 2. Now, the three conditions I mentioned can be summarized as follows.

In order to have an intentional state with content *p* a system S must:

i) have an epistemic window in which the elements that form the content *p* fall;

ii) support the attribution of another content *q* that satisfy condition *i*) and is extensionally equivalent to *p*;

iii) give rise to differential behaviors for *p* and *q* (that is: attributing *p* does not imply attributing *q*).

Now, what kind of content should be considered in this case? Is it necessary to consider just type 3 content or will type 2 content also suffice? As we have pointed out, contents of type 2 are not as fine-grained as content of type 3, but are fine-grained enough to sever the connection between truth-condition and the correctness conditions for the attribution of the associated functional states without regard to possible malfunction. I think they are sufficient in allowing the fixation of the property of *appearing*, the one that is in play when Ulysses appears as a vagrant. So, we should show that it is possible to give an interpretation of a case of animal deception that meets the conditions met by the situation in which we have supposed Ulysses and his enemies. This is the challenge.

Consider the following case, exhibited by two home-reared chimpanzees, Austin and Sherman. Austin is smaller than Sherman and subordinate to him but, unlike Sherman, he is not afraid of the dark. During the night they both rest in a hut. Because Sherman was bullying Austin throughout the day, as night approaches Austin performs a deception to reverse the dominance order. Unobserved by Sherman, he goes out of their hut, makes strange noises as though someone is scraping on the hut, then returns inside and looks outside with a worried stance. Sherman becomes fearful and stops bullying him (Savage-Rumbaugh and McDonald, 1988).[23]

This case of deception is voluntary, in that Austin would not have performed that behavior as a result of some automatic mechanism, nor if Sherman had not been nearby; and it is goal-directed, in that the aim is to induce Sherman to stop bullying him. But what is important for our purposes is that this could be considered a case in which we may interpret an animal's functional state in terms of content and then ascribe to the animal an intentional state. Let us see this conclusion in details.

May we attribute to Sherman a state whose content is: "something is making noises outside the hut"? I think we can respond affirmatively. We can attribute this state to Sherman because i) Sherman discriminates all the relevant elements of the content described and others such as noises, producing noises, and the properties of *being inside* or *outside the hut*; ii) it would be possible to interpret Sherman's internal functional state as having the content "there is my-hut-companion making noises outside the hut", an interpretation that is contingently equivalent to the state of the supposed intentional state and whose element can be discriminated by Sherman; iii) the attribution of a functional state interpreted as having the content "something is making noises outside the hut" does not imply the attribution of a functional state interpreted as having the content "there is my-hut-companion making noises outside the hut." From all this follows that Sherman can be credited with the attribution of intentional states of type 2 content that "something is making noises outside the hut".[24] Now, the attribution is secured by the epistemic window of the system, namely, by his discriminatory abilities. These abilities allow Sherman to entertain two modes of presentation related to the same truth-conditions. Hence, we have an intentional state in a non-speaking creature based on two type 2 contents. This case shows that language is not necessary in order to have intentional states.

One may object that there is no reason to attribute such complex thoughts to Sherman and Austin. It would be sufficient to consider, behavioristically, their reactions to stimuli occurring from moment to moment. This very common line of reaction, however, is somewhat paradoxical. In fact, the more the animal behavior is complex, both as a result of carefully stated anecdotes and subtle experimental settings, the more the behavioral explanations have to be articulate. The result of having very complex behavioral explanation, though, is that of making them more powerful, apt to explain also more and more cases of human behavior. So, we obtain animal non-mentality at the price of losing our own! Moreover, abandoning intentional explanation in case of mute animals reduces simplicity and predictability in ethology; and with them the advantage of ontological parsimony. Here, it seems, empirical and philosophical interests diverge.

This analysis allows us to meet such other requirements as those invoked, for instance, by Quine, Dennett and Davidson, regarding the attribution of intentional states. One of these is that intentional systems cannot exist in isolation. On the present analysis, this condition is satisfied, even if minimally, by any two speechless animals. In fact, the behavior I have discussed relies on forms of communicative interaction: the conditions for believing are set by an animal (Austin) while the belief is attributed to another (Sherman). In a certain sense, the response behavior is shared both by the deceiver and the deceived, and this lends support to the hypothesis that these animals have a natural "theory of mind" that is shared and used for communicative ends.[25] Now, since I have not presumed any special ability in the two chimpanzees, we may suppose their mental equipment is present in other individuals of their species as well.

There is also a second requirement satisfied by this approach: it is implausible to conceive an intentional system with the capacity for a single, isolated belief, because of the holistic nature of these states. My proposal, however, is set on the assumption that there should be *at least* two functional states with an equivalent content. However, such an oversimplified system would be an extreme case. For each belief we must admit that the discriminated elements of each state can be attributed as parts of other complex functional states.

5.CONCLUSIONS

As far as animals are concerned, my contention has been that we are justified in attributing intentional states to a system whenever that system exhibits:

- a behavior governed by functional states;
- the interpretation of those functional states can be put in terms of contents that are supported by discriminations that fall within its epistemic window; and
- the system exhibits behavioral patterns that satisfy certain conditions.

I have used examples of primates because they already give us plentiful evidence for grounding such attributions; but it is possible that they are not the only non-speaking

intentional systems. In looking for other possible candidates, further empirical and experimental data must be considered.

I want to conclude with two general considerations. First, since the so-called "linguistic turn" there has been a tendency to give an account of thought on the model of language. As a consequence, the possibility of any form of non-linguistic or pre-linguistic thought has been almost ruled out *a priori*. The fact is that those who have followed the "linguistic turn" have not only tended to view language as a useful methodological tool for analysis, but also as imposing a substantive constraint on the nature of thought. Consider again the problem of intentionality. If we analyze intentionality in terms of language, two main routes may be followed: either we consider linguistic expressions as the only relevant data for investigation, thereby taking language as constraining thought; or we extend the domain of the intentional to include those behaviors that satisfy certain criteria of complexity extrapolated from language. The first tradition, perhaps dominant, thinks that language mirrors thought and even if "[l]anguage may be a distorting mirror [...] it is the only mirror we have" (Dummett 1988, 7). The other tradition considers language as a means to exhibit or represent certain underlying phenomena, without taking this cognitive mechanism as necessarily linguistic in nature: in Quine's words, "Taking the objects of propositional attitudes as sentences does not require the subject to speak the language of the object sentence, or any" (1960, 213). In this chapter I have taken this second route. Our ascriptions of thoughts does not imply the ascription of intrinsically linguistic states, for such ascriptions may well be justified by the relevantly complex non-linguistic behaviors.

Second, by ascribing thoughts to animals we are not supposing that they are capable of being in exactly the same kinds of epistemic states as humans. Language makes available epistemic attitudes of greater complexity which arguably require a different level of analysis. To this end, some philosophers have contrasted beliefs with different kinds of doxastic states.[26] But if we confine ourselves to certain kinds of perceptually fixed beliefs, the cognitive capacities exercised by humans in certain situations seem to be much closer in kind to those exhibited by some speechless creatures in similar situations than they are to those exercised by humans in linguistic tasks. For all these reasons, I think that the analysis here presented may allow us to consider, from a different perspective, the relationship between thought and language.[27]

NOTES

Thorough this paper I will use "animals" to refer to living beings with behavioral capacity complex enough to raise the issue of their mentality.

See, in particular, the V section of his *Discourse on Method*, in Descartes (1637/1982).

Agar (1993); Allen (1992); Allen and Bekoff (1997); Bennett (1976); Bermudez (2003); Dennett (1983); Hurley (2003; forthcoming) Malcolm (1972); Routley (1981).

Byrne and Whiten (1988); Cheney and Seyfarth (1990); Premack and Woodruff (1978).

Bennett (1991); Davidson (1975, 1985, 1997); Dennett (1995, 1996); Heil (1992); Heyes (1993, 1998); Lowe (2000); O'Leary-Howthorne (1993); Premack (1988); Stich (1979).

Glock (2000); Peacock (1992).

Stephan (1999).

"Semantic opacity", widely regarded as the hallmark of intentionality, is a feature shared by sentences containing verbs like "believe", "desire", etc. In these sentences, the substitution of coreferential expressions may change their truth value—contrary to so-called "Leibniz's law". Related features of such sentences are their failure to satisfy the law of existential generalization and the principle of truth functionality.

Searle considers intentional states as directly referring to their truth conditions, so allowing any kind of coreferential substitution. But on such a view we would not learn anything from discovering that the Morning star *is* the Evening star, as in fact we do.

For instance: "[…] a person cannot just believe that he or she is seeing a cat; in order to believe this, one must know what a cat is, what seeing is, and above all, one must recognize the possibility, however remote, that one may be wrong" (Davidson 1999, 8).

Recognition seems to constitute an overlapping area. Some forms of recognition involve no conceptual knowledge, others do. In the following I will consider nonconceptual recognition.

E.g., McDowell (1994); Brewer (1999).

This proposal has something in common with the *Generality Constraint* by Evans; it parts company in that it allows the possibility of nonconceptual elements and in not having compositionality as a primary goal.

On this, see Allen (1992) and Allen and Bekoff (1997).

It should be emphasized that semantic opacity is not a necessary and sufficient condition for individuating intentional states. Modal talk, and even some scientific explanations, are semantically opaque.

For opposite views in appreciating this point see Davidson (1999) and Allen (1999).

The contrast between horizontal and vertical links can be also put in terms of the purposes to which attribution is put, whether to explain behavior or to facilitate communication about the world. For a unitary view of content attribution, focusing more on the first purpose, see Bilgrami (1992) and Pereboom (1995).

What follow may remind you of Dretske's analysis (1988). However, I part company from Dretske (see n. 21).

The same applies in case of the individuation of types of individuals.

See Stalnaker (1984).

Dretske (1988, 70-74) argues against this view.

For opposing view on concepts in animals, see Allen and Hauser (1991) and Chater and Heyes (1994).

Prior to this event, Sherman and Austin had been introduced to a "bad monster," that is, a person dressed in a King Kong suit, who frightened them indoors. However, I do not think that this changes the conceptual point I am making.

Notice, further, that Sherman cannot be said to believe, of something (or someone) specific, that it is making the noises outside the hut, for he has no idea what or who is making the noises. We then also have a case of failure of existential generalization across a belief context.

See Premack and Woodruff (1978); Whiten (1991), and, for a critical view, Heyes (1998).

Cohen (1992); Dennett (1978); Stich (1978).

This paper originated while I was Visiting Fellow at the Center for Cognitive Studies, Tufts University, and has been developed when I was Visiting Fellow at the Center for Cognitive Science, Rutgers University (RuCCS). I wish to thank both Centers. I am especially grateful to Dan Dennett and Brian Loar who provided constant stimulation. Thanks to Francesco Ferretti, Ausonio Marras, Carlo Penco, Silvano Tagliagambe for helpful comments on previous drafts.

Please, add these three items in bibliography, being mentioned in the new note 2 and in the modified n. 3 (s.g.)

R. Descartes (1637/1982) *Discourse on* Method, now in The *Philosophical Works of* ***Descartes***, eds E. Haldane and G. Ross, 2 vols. (Cambridge: Cambridge University Press, 1982), Vol. I,

S. Hurley (2003) "Animal Action in the Space of Reasons", *Mind and Language* 18(3): 231-256

S. Hurley (forthcoming) "Making Sense of Animals" in *Rational Animals?*, Susan Hurley and Matthew Nudds, eds., OUP, 2006.