



Wittgenstein running: Neural mechanisms of collective intentionality and we-mode

Cristina Becchio^{a,*} and Cesare Bertone^b

^a *Centre for Cognitive Science, University and Polytechnic of Turin, Via Po 14—10123 Torino, Italy*

^b *Centre for Theoretical and Applied Ontology, University of Turin, Italy*

Received 22 November 2002

Abstract

In this paper we discuss the problem of the neural conditions of shared attitudes and intentions: which neural mechanisms underlie “we-mode” processes or serve as precursors to such processes? Neurophysiological and neuropsychological evidence suggests that in different areas of the brain neural representations are shared by several individuals. This situation, on the one hand, creates a potential problem for correct attribution. On the other hand, it may provide the conditions for shared attitudes and intentions.

© 2003 Elsevier Inc. All rights reserved.

Keywords: Collective intentionality; We-mode; Mirror neurons

1. Introduction

Once Wittgenstein went to the railway station with a friend, who had to take a train. It was late and when they arrived in the vicinity of the station, they began to run. Wittgenstein was faster than his friend and reached the platform before her. At that point, unexpectedly, he jumped onto the departing train. His friend remained on the platform astonished, while the train moved away.

This anecdote, recounted by Johnson-Laird during a lesson—the truth of which we would not be prepared to swear to—describes a case of what might be called “intentional contagion.” Wittgenstein did not intend to take the train. His action was contaminated by his friend’s intention.

But how is it possible to be contaminated by the intentions of others?

The solution to this riddle may be provided by an unlikely source: the study of single neurons in the brains of monkeys.

* Corresponding author. Fax: +39-011-8159039.

E-mail address: becchio@psych.unito.it (C. Becchio).

Giacomo Rizzolatti and his group have found that certain cells located in the ventral premotor cortex (area F5) of the monkey's brain discharge during active movements of the hand or mouth, or both. In most F5 neurons, this discharge correlates with action, rather than with individual movements constituting that action. Grasping, manipulating, holding objects are the most effective actions triggering a response. One might be tempted to think that these are motor "command" neurons, making muscles do certain things. However, the astonishing truth is that these neurons also fire when the monkey in question observes another monkey—or even the experimenter—performing the same action (Gallese, Fadiga, Fogassi, & Rizzolatti, 1996; Rizzolatti, Fadiga, Gallese, & Fogassi, 1996). Any given mirror neuron responds both when a particular action is performed by the recorded monkey and when a similar gesture is performed by another individual. Neither the sight of the object itself nor the agent itself are sufficient to evoke a neural response: in order to be triggered, mirror neurons require a specific observed action.

The majority of mirror neurons respond selectively when the monkey observes a given type of action, showing congruence between the observed and executed action. This congruence may be almost perfect: the motor action may correspond to the action that, when seen, triggers the neuron.

Mirror neurons were discovered in monkeys but we know from neuropsychological studies that they also exist in the human brain.

In an interesting study Fadiga, Fogassi, Pavesi, and Rizzolatti (1995) tested the excitability of the motor cortex of normal human subjects by using transcranial magnetic stimulation (TMS), a technique used to stimulate relatively restricted areas of the cerebral cortex. The basic assumption underlying the experiment was the following. If the observation of actions activates the premotor cortex in humans, as it does in monkeys, this mirror effect should elicit an enhancement of the motor evoked potentials induced by transcranial magnetic stimulation of the motor cortex. Motor evoked potentials (MEPs) were recorded from various arm and hand muscles. The results showed that during the observation of grasping, as predicted, MEPs increased significantly. Even more intriguing was the finding that the increase of excitability was present only in those muscles that subjects would use when actively performing the observed action.

These results pose the problem of the anatomical localisation of the mirror system within the human brain. This issue has been addressed by brain imaging experiments utilising Positron Emission Tomography (PET) and, in more recent years, functional magnetic resonance (fMRI). Early studies showed that during the observation of different types of hand grips performed on a variety of objects there is an activation of the left STS region, of the inferior parietal lobule and of Broca's area (see Decety et al., 1997; Grafton, Arbib, Fadiga, & Rizzolatti, 1996; for review, see Allison, Puce, & McCarthy, 2000; Decety & Grèzes, 1999).

In a recent fMRI study by Buccino et al. (2001), subjects were shown object-directed actions made with the mouth (e.g., biting an apple), the hand (e.g., grasping a cup) and the foot (e.g., kicking a ball). The results showed that observation of object-related actions determines, among other activations, a somatotopically organized activation of the premotor cortex. During mouth-action observation there is a bilateral activation of the ventral premotor cortex (area 6) plus an activation of Broca's area. During hand-action observation a more dorsal part of ventral area 6 plus Broca's area are recruited. Finally, the observation of foot actions elicits an activation of dorsal area 6.

All these results suggest that we have a "mirror matching system" similar to that discovered in monkeys. Whenever we are looking at someone performing an action, in addition to the activation of various visual areas, there is a concurrent activation of motor circuits. Even though we

do not overtly reproduce the observed action, our motor system nevertheless becomes active as if we were indeed executing that very same action we are observing. This notion is corroborated by evidence coming from neurological patients.

Some of the patients affected by dementia show an impulsive tendency to imitate other people's movements. Imitation is performed immediately, with the speed of a reflex action, and concerns gestures that are commonly executed as well as those that are rare and even bizarre. It can be hypothesized that this imitation behaviour, called echopraxia, represents a release of a covert action simulation also present in normal subjects, but normally inhibited in its expression by the cortical areas that are functionally defective in demented patients (Gallese, 2001).

Given the existence of a system which matches executed and observed actions, the question is: what might the function of such a matching system be? When new motor skills are learned, one often spends the first training phases trying to replicate the movements of an observed instructor. Mirror neurons could in principle facilitate this kind of learning. Moreover, as Rizzolatti and Arbib (1998) note, these neurons may also enable one to mime—and possibly understand—the lip and tongue movements of others which, in turn, could provide the opportunity for language to evolve. Other studies explored the possibility that mirror neurons underlie the process of mind-reading, or serve at least, as a precursor to such a process (Gallese & Goldman, 1998).

In this paper we will discuss the relevance of the discovery of a mirror system with respect to another possibility: that mirror neurons are what enable us to share an intention with another.

We will begin our account by discussing the notion of collective intentionality and presenting some of the problems it entails.

2. Collective intentionality and We-mode

Jointly writing a paper, carrying a table, singing a duet, voting in elections, creating and upholding social institutions are all examples of collective behaviour. But what is the difference between collective and individual behaviour?

Imagine, suggests Searle (1995) a group of people sitting on the grass in various places in a park. Suddenly it begins to rain and they all get up and run to a common, centrally located shelter. In this situation each person has the intention 'I am running to the shelter,' independently of one another. The actions are individual and do not involve anything similar to collective behaviour. Imagine now a situation in which there is a group of theatre actors performing a play in which everybody performs the same action as in the earlier case, only now as part of their collective performance. What makes these cases different? Since the body movements are exactly the same, the difference must lie in the mental component. The latter involves, according to Searle, a collective intention of the form 'we intend to do x,' a collective intention which is different from an individual intention of the form 'I intend to do x.'¹

Tuomela (2001) describes collective intentions as "we-mode" intentions. The notion of "we-mode" depends on the notion of functioning (thinking and acting) as a group member versus

¹ Searle's example focuses on the difference between collective intentionality and individual intentionality. What remains identical in the two situations is the environmental context: the actions, whether motivated by an individual or collective intention, take place in a context which is shared. A framework for integrating different structures of sharedness is presented in Becchio and Bertone (2002).

functioning as a private person. If I have an attitude or act as a private person then I have an attitude or act in an “I-mode.” If I have an attitude or act *qua* a group member then I have a “we-mode” attitude or act in “we-mode.”

But what does acting as a group member mean? Basically, it involves acting intentionally within the realm of group concern. Such actions can be either successful or unsuccessful. What is required, notes Tuomela, is that the group member in question intentionally attempts to act in a way related to what he takes to be the realm of the group’s concerns in such a way that he does not violate the group’s central, constitutive goals, standards and norms, that is to say, the ethos of the group. In contrast to functioning as a group member, one can be a group member without always acting as a group member.

The literature on collective intentionality includes many dilemmas and paradoxes. Among them, the problem of *Friday’s Footprint* may be useful to bring to the fore the conditions essential for shared intentions. The problem is the following: was Robinson Crusoe truly isolated before he saw Friday’s footprint in the sand? (Brothers, 1997). In terms of intentionality, could Robinson Crusoe have had collective intentions before meeting Friday?

According to Searle, mental states exist in the brain and the same can be said for intentionality. Collective intentions also exist entirely in the heads of individuals. What makes we-intentions special is not the type of entity we-intentions can be attributed to, but the psychological mode which characterizes them: “all the intentionality needed for collective behaviour can be possessed by individual agents even though the intentionality in question makes reference to the collective” (Searle, 1990). In this view, the idea that a single individual has a collective intention appears plausible: even before meeting Friday, Robinson Crusoe could have had collective intentions.

The idea that a single individual can have a collective intention also seems possible in Tuomela’s account of collective intentionality. If the condition for we-mode is to function, thinking and acting *qua* a group member, which means not violating the ethos of the group, we can certainly admit that Robinson Crusoe could have had collective intentions even before meeting Friday. Although isolated, he continued in fact to live in observance of the standards and norms of English society, that is, of the group of which he was a member.²

It should be noted that the claim that all collective intentionality is in the individual mind or brain, does not imply that collective intentions are reducible to individual intentions. This consideration introduces what has been called the “irreducibility problem:” are we-intentions reducible to I-intentions?

According to Searle (1995, 1998), collective intentionality cannot be analysed as the sum of individual behaviour or intention. Individual intentionality and collective intentionality are

² Nevertheless, seeing a footprint is for Robinson a thunderstruck experience. “But now I come to a new scene of my life. It happened one day, about noon, going towards my boat, I was exceedingly surprised with the print of a man’s naked foot on the shore, which was very plain to be seen on the sand. I stood like one thunderstruck, or as if I had seen an apparition. I listened, I looked round me, but I could hear nothing, nor see anything; I went up to a rising ground to look farther; I went up the shore and down the shore, but it was all one; I could see no other impression but that one. I went to it again to see if there were any more, and to observe if it might not be my fancy; but there was no room for that, for there was exactly the print of a foot toes, heel, and every part of a foot.” (From: Daniel Dafoe, *The life and strange surprising adventures of Robinson Crusoe*, 1719). A footprint in the sand means the possibility of another individual, possibly a cannibalistic savage and thus, of another I irreducible to the we-dimension to which Robinson belongs. The print of an English boot would have made a completely different impression.

different and equally primitive kinds of biological phenomena. Neither of them can be reduced to the other.

Tuomela (2003) also regards we-intention as an irreducible phenomenon: we-mode represents group-level thinking and acting and is not reducible to I-mode, nor vice versa.

The irreducibility of collective intentions is a highly controversial point. It may in fact be objected that, although collective intentional behaviour is not analysable as simply the sum of individual intentional behaviour, collective intentionality is not a primitive phenomenon, but a sequence of individual intentions related in a certain way: if “we intend”, then each of us has an intention of the “I intend” form, plus a belief about the other’s intention. In this manner, collective intentionality is reduced to individual intention plus a belief about the other’s intention (Tuomela & Miller, 1988).

Bratman (1992) agrees with Searle that collective intentional behaviour is not analysable as just the summation of individual intentional behaviour even supplemented by mutual beliefs. Yet contrary to Searle, he thinks that collective intentionality is not a primitive phenomenon and retains that a reductive analysis may be possible. The conditions essential for shared cooperative activity are summarized by Bratman (1992) in the following way:

Where J is a cooperatively neutral joint-act type, our J-ing is shared cooperative activity only if:

- (1) (a) (i) I intend that we J
- (1) (a) (ii) I intend that we J in accordance with and because of meshing subplans of (1) (a) (i) and (1) (b) (i)
- (1) (b) (i) You intend that we J
- (1) (b) (ii) You intend that we J in accordance with and because of meshing subplans of (1) (a) (i) and (1) (b) (i)
- (1) (c) The intention in (1) (a) and (1) (b) are minimally cooperatively stable
- (2) It is common knowledge between us that (1)

The originality of Bratman’s analysis comes from conditions (1) (a) (ii) and (1) (b) (ii), which are meant to capture the feature of mutual responsiveness. For an activity to be a shared cooperative activity, each agent must intend that the group performs this action in accordance with subplans that mesh. This meshing of subplans in turn implies that the intentions of the participants must be interlocked. By thus requiring, Bratman moves away from the classical reductive analyses of collective intention according to which the crucial link among the attitudes of the participating agents is simply cognitive.

One drawback to Bratman’s analysis is that it seems tailored to a form of shared cooperative activities that are already rather sophisticated. Bratman’s analysis focuses on future-directed intentions, that involve rational deliberation and planning. One may therefore wonder whether it still applies when one considers more elementary forms of collective behaviour.

Furthermore, condition (1) (a) (i) presupposes that a ‘we’ has already been established: “I intend that we J.” But how can I intend that ‘we’? How is this ‘we’ possible? Where is it located?

On our view, the possibility to enter in relation with other individuals so as to constitute a ‘we,’ presupposes a collective intentionality. Thus, in default of a primitive ‘we,’ the condition cannot be “I intend that we J,” but “I intend that I,” “You intend that you”.

The critique that in general can be raised against reductive approaches is that they are cognitively implausible: in order to conclude that a given intention or belief is definitely common, all the interlocutors should commit themselves to carrying out a series of inferences set out in every

time. Even if we conceded that we stop after a reasonably short time, for example at the third or fourth level, sharedness would become a strenuous undertaking for adults, and virtually impossible for children under the age of ten: there are too many embeddings, and these require cognitive resources unavailable under a certain age (Clark, 1996; see also Bara, 1999).

In the next section we will consider the problem of the neural conditions of shared attitudes, that is, which neural mechanisms underlie “we-mode” processes or serve as precursors to such processes?

3. Neural basis of We-mode

As suggested by Gallese (2001), mirror neurons may constitute the neural substratum of a “resonance mechanism” which enables us to resonate with others. Consider the properties of mirror neurons. As we have seen, the very same mirror neuron fires not only when the monkey executes an action, but also when the immobile monkey in question observes another monkey performing the same action. What does this indicate? In terms of brain representation, what is the possible meaning underlying this behaviour?

According to Gallagher and Jeannerod (2002), it indicates that in at least one point of the brain, representation for action is shared among individuals. This discovery has the astonishing effect of reversing the problem, which is no longer “how is a shared intention possible?”, but “how can one distinguish one’s action from those of other people?” Paradoxically, we move from the problem of recognizing the others’ intentions and we find ourselves faced with the opposite problem: distinguishing self-produced actions from actions performed by others.

Georgieff and Jeannerod (1998) define this the question of “Who?” Brain mapping experiments (using PET or fMRI) show activation of a partially overlapping cortical and subcortical network during motor imagery and action observation. This network involves structures directly concerned with motor execution, such as motor cortex, dorsal and ventral premotor cortex, lateral cerebellum, basal ganglia; it also involves areas concerned with action planning, such as dorsolateral prefrontal cortex and posterior parietal cortex. A recent meta-analysis of these data (Grèzes & Decety, 2001) reveals that the degree of overlap between different modalities of representations varies from one cortical area to another. With regard to primary motor cortex itself, fMRI studies demonstrate that pixels activated during contraction of a muscle are also activated when a movement involving the same muscle is imagined (e.g., Roth et al., 1996). Electrophysiological measurements (Hari et al., 1998) have also shown that when subjects observe hand movements, there is a desynchronization over the motor cortex similar to that which occurs during actual movements. In premotor cortex and SMA, the overlap between imagined and observed actions is almost complete. By contrast, action observation largely involves inferotemporal cortex, which is not the case for action imagination.

Several imaging studies of imitations (Chaminade, Meltzoff, & Decety, 2002; Decety, Chaminade, Grèzes, & Meltzoff, 2002; Iacoboni et al., 1999) show an activation of the parietal lobe during both action execution and action observation. Importantly, these studies suggest that the parietal cortex plays a crucial role not only when the self-resonates with others, but also in distinguishing the self from the other.

According to Gallagher and Jeannerod (2002), this relative similarity of neurophysiological mechanisms accounts both for the fact that actions can normally be attributed to their real author, and that action attribution remains a fragile process.

Diagram in Fig. 1 helps illustrate why.

The diagram represents a motor cognitive situation with two people. We have agent A and agent B. The processes diagrammed are represented as occurring in the brain. They are based on the idea of a partial overlap between neural representations that are constructed when an action is observed or when an action is thought of. Let us make an example: agent A generates a representation of a self-generated action, a motor intention. If this is then executed, it will become a signal for agent B, such that agent B will form a representation of the action that he sees. Agent B will make an estimate of the social consequences of the action he sees and will possibly change his beliefs about agent A. The interaction of the two agents will depend on the interaction of the representations of the observed and executed actions inside each of the two brains. In fact the two representations, within an individual subject, are close to each other and partly overlap. Determining who is acting, myself or the other, will be based on the non-overlapping part.

The potential problem in correctly attributing an action to its agent was systematically examined experimentally (Daprati et al., 1997). The Subject's hand and the Experimenter's hand were filmed by two different cameras. By changing the position of a switch, one or the other hand could be briefly

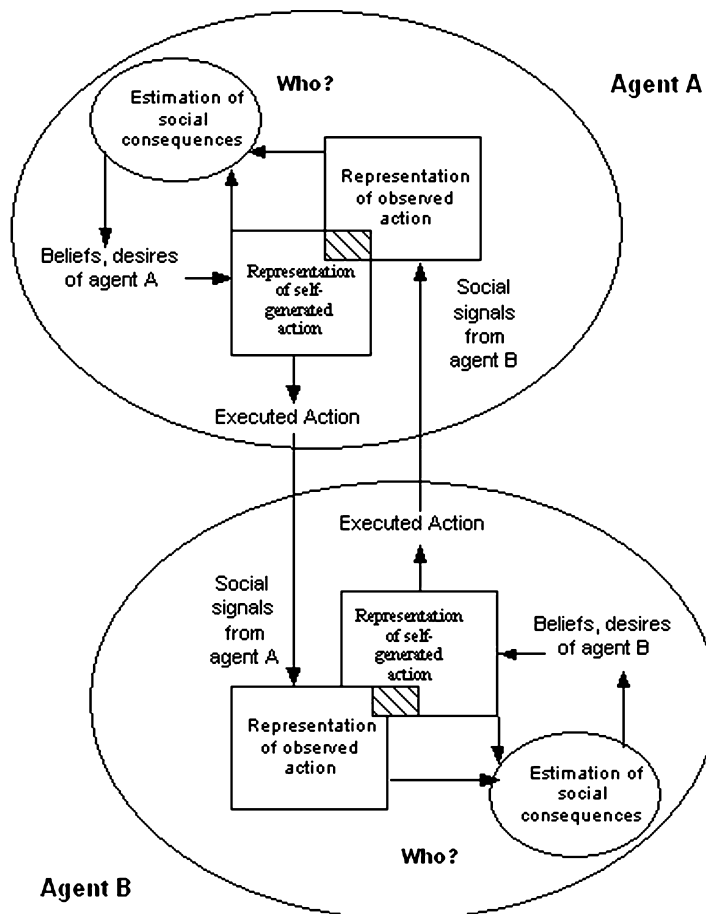


Fig. 1. From Gallagher and Jeannerod (2002).

displayed on the video screen seen by the Subject. The two hands looked alike as they were covered with a similar glove. In each trial, both the Experimenter and the Subject had to perform a given hand movement on command (e.g., stretch thumb, stretch fingers 1 and 2, etc): on some trials, however, the Experimenter's movement departed from the instruction. As a result of this experimental arrangement, the subject was randomly shown either his own hand, or the Experimenter's hand performing the same movement as his, or a different movement. At the end of each trial, a verbal judgement was recorded: the subject had to say whether the hand he had seen was his own hand or another hand. Normal subjects were able to unambiguously determine whether the moving hand seen on the screen was theirs or not, in the two "easy" conditions: First, when they saw their own hand, they correctly attributed the movement to themselves; second, when they saw the experimenter's hand performing a movement which departed from the instruction they had received, they correctly denied seeing their own hand. By contrast, their performance worsened in the "difficult" trials where they saw the Experimenter's hand performing the same movement as required by the instruction: in this condition, they misjudged the alien hand as their own in about 30% of cases.

In patients presenting delusions, this difficulty dramatically increases and actions become systematically misattributed. Using the same paradigm as above in groups of schizophrenic patients, Daprati et al. (1997) found a dramatic increase in the rate of incorrect responses in the "difficult" trials. The error rate was 80% in a group of schizophrenics with delusional experiences, whereas in a non-hallucinating group, it was only 50%. The fact that all patients gave nearly correct responses in the other two conditions (the error rate remained within 1–7%) excludes the possibility that the effect observed in the "difficult" trials could be due to factors unrelated to the task, such as lack of attention.

4. Sharing sensations and emotions

So far we have provided evidence supporting the notion that there are areas of the brain where the representation of *action* is shared among individuals. But what about sensations, pains and emotions? Could a mirror matching system also exist for pains, emotions and sensations?

This hypothesis can be supported by empirical data. Preliminary evidence suggests that in humans a "mirror phenomenon" occurs in pain related neurons. Hutchinson, Davis, Lozano, Tasker, and Dostrovsky (1999) studied pain related neurons in the human cingulate cortex. Cingulotomy procedures for the treatment of psychiatric illnesses provided an opportunity to examine the responsiveness of neurons located in the anterior cingulate cortex to painful stimuli prior to excision. It was noticed that a neuron that responded to noxious mechanical stimulation applied to the patient's hand also responded when the patient watched pinpricks being applied to the examiner's fingers. Both applied and observed painful stimuli elicited the same response in the same neuron.

A recent study published by Calder, Keane, Manes, Antoun, and Young (2000) examined a Huntington patient selectively impaired in detecting disgust in different modalities, such as facial signals, emotional sounds, and emotional prosody. As one would expect in the hypothesis of damage to the emotional mirror system, the same patient was also selectively impaired in subjectively experiencing disgust and therefore in reacting appropriately to it.

What this preliminary evidence indicates is that the same structures that are active during sensations and emotions are also active when the same sensations and emotions are detected in others. It appears therefore that a whole range of "mirror matching systems" may be present in

the human brain. In other words, “mirror phenomena” are not to be seen as limited to a particular group of motor neurons in the ventral cortex, but as a modality of functioning which is widespread in the brain.

5. Back to the problem of shared intentions

We are now in a position to deal with the problem from which we started: how is shared intentionality possible?

The empirical data that we have reviewed suggest that in different areas of the brain (overlapping areas in Fig. 1) neural representations underlying action observation and execution *share* a basic neural format.

Consider once again the properties exhibited by motor mirror neurons. As we have seen, the same neuron discharges both when the monkey executes an action and when a similar action is performed by an other individual. This requires that the neuron discharge is neither purely visual nor purely motor: motor and sensory properties must coexist within the same neuron. Multi-modality is nevertheless a necessary, but not a sufficient condition. If the executed action has to overlap the observed action, the representation has to be not only multi-modal, but multi-subjective. The action cannot be mapped in the perspective of the executor, nor in that of the observer: the format has to be such that the action can be referred to both the executor and the observer. We propose to define “we-centric”³ a neural format which exhibits these characteristics.

We-centric representations constitute the neural basis of the social mind. It is by means of we-centric representations that we recognize other humans beings as similar to us and are able to mirror ourselves in their behaviour. In absence of mirror resonance mechanisms, witnessing the action of others would not be different from seeing any other physical movement.

Sharing joint intentions depends, at a basic level, on the same mechanisms. Two agents acting collectively the same action share a we-centric representation of the action in question. The neural representation of the executed action tends to overlap the representation of the observed action: as a result, the performed action can be indifferently ascribed to the self or the other.⁴

³ In Becchio and Bertone (2002) we define “we-centric” the neural format of intersubjectivity. The term “we-centric” appears with this sense in Gallese (2003).

⁴ On this point our analysis draws away from that by Gallese (2001). Starting from a neurobiological standpoint (the discovery of a mirror matching system), Gallese (2001) analyses how actions are possibly represented and understood. His hypothesis is that action recognition implies an implicit process of action simulation mediated by mirror matching systems. On this hypothesis, “much of what we ascribe to the mind of others when witnessing their actions depends on the “resonance mechanisms” that they trigger in us.” “It is possible to speculate,” continues Gallese, “that this system may originally have developed to achieve a better control of action performance. The coupling of this forward model architecture with the vision of the agent’s hand, and its later generalization to the hands of others, may have allowed this system to be used also for totally different purposes, namely to represent other individuals’ actions.” Following this hypothesis, intersubjectivity can be viewed as the result of a process of “motor equivalence” that equates the action observed with the executed action.

We agree with Gallese (2001) that mirror neurons function as a “resonance mechanisms.” Yet, it seems to us that the speculative scenario sketched by Gallese suffers from a subjective prejudice. Gallese ascribes to the subject a primacy that the subject, as least as regards F5, seems not to have. In other words, if we want to take the properties of mirror neurons seriously, then we need to consider that in F5 the agent’s and the observer’s representations are interchangeable and that the subject is not represented in a preferential way. For a detailed discussion, see Becchio e Bertone (2002).

Nevertheless, many forms of social activities require not only sharing—and eventually, imitating—what another individual is doing, but performing a different action as part of collective intentional behaviour. Communication is a striking example of “cooperation without imitation.” Other examples are playing chess or basket ball, acting in a play, taking a seat at the cinema. In all these cases, it would be harmful to simply replicate the actions of another: what is required is exactly the opposite, that is, to avoid doing what others have already done. At this level of complexity, the existence of non-overlapping parts, as well the existence of possible differences in intensity of activation between activated areas, is necessary to each agent for discriminating between himself and the other, between the actions that he observes and the actions that he needs to perform.

6. Conclusions

In conclusion, at a neural level the difference between individual and collective intentions, I-mode and we-mode does not depend on contents nor on the type of entities (individuals versus collective entities), but on the format of the underlying neural representations.

Converging neuroscientific data suggest that in different areas of the brain neural representations underling the self and the other's behaviour share a common, we-centric code.

The idea that we propose in this paper is that we-centric neural representations represent the essential neural condition for collective intentional behaviour.

Acknowledgments

This research was supported by the Ministero Italiano dell'Istruzione dell'Universita' e della Ricerca (Firb Project, “Assessment dei disturbi della comunicazione in un'ottica riabilitativa” research code n. RBAU01JEYW_001).

References

- Allison, T., Puce, A., & McCarthy, G. (2000). Social perception from visual cues: Role of the STS region. *Trends in Cognitive Sciences*, 4, 267–278.
- Bara, B. G. (1999). *Pragmatica Cognitiva*. Torino: Bollati Boringhieri.
- Becchio, C., & Bertone, C. (2002). Il problema della condivisione. *Sistemi Intelligenti*, 14(2), 207–216.
- Bratman, M. E. (1992). Shared cooperative activity. *The Philosophical Review*, 101(2), 237–241.
- Brothers, L. (1997). *Fridays Footprint: How Society shapes the Human Mind*. Oxford: Oxford University Press.
- Buccino, G., Binkofski, F., Fink, G. R., Fadiga, L., Fogassi, L., & Gallese, V. (2001). Action observation activates premotor and parietal areas in a somatotopic manner: An fMRI study. *European Journal of Neurosciences*, 13, 400–404.
- Calder, A. J., Keane, J., Manes, F., Antoun, N., & Young, A. W. (2000). Impaired recognition and experience of disgust following brain injury. *Nature Neuroscience*, 3(11), 1077–1078.
- Chaminade, T., Meltzoff, A. N., & Decety, J. (2002). Does the end justify the means? A PET exploration of the mechanisms involved in human imitation. *Neuroimage*, 15(2), 318–328.
- Clark, H. (1996). *Using Language*. Cambridge, UK: Cambridge University Press.

- Daprati, E., Franck, N., Georgieff, N., Proust, J., Pacherie, E., Dalery, J., & Jeannerod, M. (1997). Looking for the agent: An investigation into consciousness of action and self-consciousness in schizophrenic patients. *Cognition*, 65, 71–86.
- Decety, J., Chaminade, T., Grezes, J., & Meltzoff, A. N. (2002). A PET exploration of the neural mechanisms involved in reciprocal imitation. *Neuroimage*, 15(1), 265–272.
- Decety, J., Grezes, J., Costes, N., Perani, D., Jeannerod, M., Procyk, E., Grassi, F., & Fazio, F. (1997). Brain activity during observation of actions. Influence of action content and subject's strategy. *Brain*, 120, 1763–1777.
- Decety, J., & Grèzes, J. (1999). Neural mechanisms subserving the perception of human actions. *Trends in Cognitive Sciences*, 3, 172–178.
- Fadiga, L., Fogassi, L., Pavesi, G., & Rizzolatti, G. (1995). Motor facilitation during action observation. A magnetic stimulation study. *Journal of Neurophysiology*, 73, 2608–2611.
- Gallagher, S., & Jeannerod, M. (2002). From action to interaction. *Journal of Consciousness Studies*, 9, 3–26.
- Gallese, V. (2001). The 'shared manifold' hypothesis. From mirror neurons to empathy. *Journal of Consciousness Studies*, 8, 33–50.
- Gallese, V., Fadiga, L., Fogassi, L., & Rizzolatti, G. (1996). Action recognition in the premotor cortex. *Brain*, 119, 593–609.
- Gallese, V., & Goldman, A. (1998). Mirror neurons and the simulation theory of mind reading. *Trends in Cognitive Science*, 2, 493–501.
- Georgieff, N., & Jeannerod, M. (1998). Beyond consciousness of external reality: A "Who" system for consciousness of action and self-consciousness. *Consciousness and Cognition*, 7, 465–477.
- Grafton, S. T., Arbib, M. A., Fadiga, L., & Rizzolatti, G. (1996). Localization of grasp representations in humans by PET: 2. Observation compared with imagination. *Experimental Brain Research*, 112, 103–111.
- Grèzes, J., & Decety, J. (2001). Functional anatomy of execution, mental simulation, and verb generation of actions: A meta-analysis. *Human Brain Mapping*, 12, 1–19.
- Hari, R., Forss, N., Avikainen, S., Kirveskari, E., Salenius, S., & Rizzolatti, G. (1998). Activation of human primary motor cortex during action observation: A neuromagnetic study. *Proceedings of the National Academy of Sciences of the United States of America*, 95(25), 15061–15065.
- Hutchinson, W. D., Davis, K. D., Lozano, A. M., Tasker, R. R., & Dostrovsky, J. O. (1999). Pain-related neurons in the human cingulate cortex. *Nature Neuroscience*, 2(5), 403–405.
- Iacoboni, M., Woods, R. P., Brass, M., Bekkering, H., Mazziotta, J. C., & Rizzolatti, G. (1999). Cortical mechanisms of human imitation. *Science*, 286(5449), 2526–2528.
- Rizzolatti, G., & Arbib, M. A. (1998). Language within our grasp. *Trends in Neurosciences*, 21(5), 188–194.
- Rizzolatti, G., Fadiga, L., Gallese, V., & Fogassi, L. (1996). Premotor cortex and the recognition of motor actions. *Cognitive Brain Research*, 3, 131–141.
- Roth, M., Decety, J., Raybaudi, M., Massarelli, R., Delon-Martin, C., Segebarth, C., Morand, S., Gemignani, A., Décorps, M., & Jeannerod, M. (1996). Possible involvement of primary motor cortex in mentally simulated movement. A functional magnetic resonance imaging study. *Neuroreport*, 7, 1280–1284.
- Searle, J. R. (1990). Collective intentions and actions. In P. R. Cohen, J. Morgan, & M. E. Pollack (Eds.), *Intentions in Communication*. Cambridge (MA): MIT Press.
- Searle, J. R. (1995). The construction of social reality. *Penguin Philosophy*.
- Searle, J. R. (1998). *Mind, Language and Society*. New York: Basic Books.
- Tuomela, R. (2003). The We-Mode and the I-Mode. Forthcoming. In F. Schmitt (Ed.), *Socializing Metaphysics: The Nature of Social Reality*. Lanham, Md: Rowman and Littlefield.
- Tuomela, R. (2001). Collective Intentionality and Social Agents. Paper invited for the AI conference IMF, Toulouse, May 21–23.
- Tuomela, R., & Miller, K. (1988). We-intentions. *Philosophical Studies*, 53, 367–389.