

**I FEEL YOUR PAIN:
MIRROR NEURONS, EMPATHY, AND MORAL MOTIVATION**

*Paul Thagard
University of Waterloo*

Thagard, P. (2007). I feel your pain: Mirror neurons, empathy, and moral motivation. *Journal of Cognitive Science*, 8, 109-136.

ABSTRACT: Mirror neurons are brain systems found in monkeys and humans that respond similarly to actions and to the perception of actions of others. This paper explores the implications of mirror neurons for several important philosophical problems, including knowledge of other minds, the nature of empathy, and moral motivation. It argues that mirror neurons provide a more direct route to other minds, empathy, and moral motivation that complements the more familiar route based on conscious, verbal inference. To show how mirror neurons accomplish these functions, I apply a neurocomputational account of representation and inference.

INTRODUCTION

The discovery of mirror neurons has been hailed as one of the major recent breakthroughs in neuroscience, with possible implications for the explanation of many important cognitive functions, including action understanding, imitation, language, and empathy. Mirror neurons were first identified in the 1990s by Giacomo Rizzolatti and his colleagues at the University of Parma (Rizzolatti and Craighero, 2004). They found that monkey prefrontal cortex contains a particular class of neurons that discharge both when the monkey does a particular action and when it observes another individual doing a particular action. Similar classes of neurons have been found in humans, capable of mirroring not only physical actions but also pain and disgust.

September 22, 2009

This paper is an exploration of the epistemological and ethical significance of mirror neuron mechanisms. After briefly reviewing the recent neuroscience literature on mirror neurons, I will discuss its implications for the traditional philosophical problem of other minds. Mirror neurons seem to offer a more direct route to the understanding of other people than provided by the usual kinds of inference – analogical and explanatory – that have been considered to provide solutions to the problem of other minds. Similarly, mirror neurons shed light on the nature of empathy, with implications for the moral psychology of caring. I argue that mirror neurons help with the problem of moral motivation, which concerns the relation between moral judgments and people’s willingness to act on them. Finally, I use a neurocomputational account of representation and inference to explicate the kinds of understanding that mirror neurons help to provide.

MIRROR NEURONS

When a monkey grasps an object, there are neurons in area F5 of its premotor cortex that fire. Much more surprising is the serendipitous discovery by Rizzolatti and his colleagues that the same region contains neurons that fire both when the monkey grasps an object and when it observes another monkey or human grasping an object (Di Pellegrino et al., 1992; Gallese et al., 1996; Rizzolatti et al., 1996; Nelissen et al., 2005). There are mirror neurons in F5 for grasping both with hands and with mouths, and another area, the superior temporal sulcus, contains mirror neurons for walking, turning the head, bending the torso, and moving the arms. The observations represented by mirror neurons are visual-motor, integrating the visual and motor experiences of monkeys. Rizzolatti and Craighero (2004) argue that the mirror neuron system is the

basis for both action understanding and imitation. Not only does a monkey's mirror neuron system give it a direct appreciation of what another monkey is doing when it moves, but also facilitates imitating those motions that might be useful for its own goals such as finding food. Mirror neurons can also work with auditory-motor representations: Kohler et al. (2002) found neurons in monkey premotor cortex that discharge when the animal performs a specific action and when it hears the related sound.

The evidence for mirror neurons in monkeys comes from direct recording of single neurons, but evidence for analogous systems in humans is largely indirect, from brain scanning. Rizzolatti and Craighero (2004) cite many studies that show that the observation of actions done by others activates in humans a complex network formed by visual and motor areas. According to Rizzolatti (2005a), evidence that a mirror system exists in humans comes from electroencephalography, magnetoencephalography, transcranial magnetic stimulation, and brain imaging studies. Hence observing the physical actions of others prepares people not only to understand what they are doing but also to imitate them. More controversially, Rizzolatti and Arbib (1999) conjecture that mirror-neuron systems provided the neurophysiological basis from which language developed as an extension of gestural communication. Iacoboni et al. (2005) argue that premotor mirror neurons are also involved in understanding others' intentions. Later I will try to give an account of the nature of such understanding.

In humans, mirror neurons may be relevant for how people understand emotions as well as actions. Gallese, Keysers, and Rizzolatti (2004) claim that a mirror-neuron system involving visceral-motor centers enables people to understand each other's emotions, just as one involving visual-motor centers enables people to understand each

other's actions. Wicker et al. (2003) used fMRI brain scans to compare how people react to disgusting smells with how they react to video clips of people reacting to disgusting smells. They found that the brain's anterior insula, which is known to collect information from various visceral centers, is activated both during the emotion of disgust evoked by unpleasant odorants and during the observation of facial expressions of disgust. Additional overlap was found in the anterior cingulate cortex. Hence it appears that these two cortical areas, the insula and the anterior cingulate, enable people to appreciate other peoples' emotions of disgust.

Similarly, Botvinick et al. (2005) used neuroimaging to find that perception of facial expressions of pain engage cortical areas also engaged by the first-hand experience of pain, including the anterior cingulate and the insula, the same areas that had been found to mirror disgust. Singer et al. (2004, 2006) also found that insular and anterior cingulate cortex were activated both by receiving pain and by observing a loved one receiving pain. Further support for the mirroring of pain is found in the studies of Avenanti et al. (2005) who used transcranial magnetic stimulation to detect evidence for the presence of empathic appreciation of the sensory qualities of the pain of others. I shall discuss later the significance of mirror neurons and empathy and mirror neurons for moral psychology. First, however, I want to explore the relevance of mirror neurons to the traditional philosophical problem of other minds.

OTHER MINDS

There are at least two problems of other minds. The first, philosophically skeptical one, concerns how you can be justified at all in believing that there are other minds besides your own. You know that you have a mind because you directly

experience your own thoughts, perceptions, and emotions, but you have no direct experience of the mental states of others. So how are you justified in believing that other people have mental states?

The second, more moderate problem of other minds assumes that we can have knowledge about the minds of others, but asks how we manage to do this. What kinds of reasoning or experience enable us indirectly to know what is going on in the minds of people around us? This moderate problem merges with the psychological problem of the nature of the cognitive processes that enable people to understand each other: it belongs as much to the branch of social psychology called social cognition as it does to the philosophy of mind. This section argues that mirror neurons help with both the skeptical and the social-cognition problems of other minds.

Traditional philosophical solutions to the skeptical problem attempt to identify the kinds of inference that justify belief in other minds (Hyslop, 2005). John Stuart Mill and others have construed the inference to other minds as analogical, with a structure such as the following:

I know by introspection that I have beliefs, desires, and emotions that produce my range of behaviors.

Other people are very similar to me in that they have the same range of verbal and physical behaviors.

Therefore, by analogy, other people also have mental states.

Unfortunately, this argument is rather weak, since it does not deal with the fundamental difference between myself and others, namely that I directly experience my own mental states but simply cannot do so for others. The argument also fails to take into account

possible alternative explanations of why other people behave like me, for example that they are robots or zombies or remotely controlled, simulating mental states without actually having them.

More plausibly, belief in the existence of other minds can be seen as justified by a kind of reasoning called *inference to the best explanation* (Graham, 1998). This kind of inference, which is common in science as well as everyday life, consists in accepting hypotheses if they provide a better explanation of all the available evidence than competing hypotheses. For example, a detective might conclude that one suspect is a murderer because that hypothesis is the best explanation of the forensic evidence such as the presence of the suspect's fingerprints on the murder weapon. For the problem of other minds, the inference runs as follows:

Other people have a wide range of verbal and physical behaviors.

One hypothesis that would explain their behaviors is that they have mental states.

This hypothesis provides a better explanation of the evidence than alternative hypotheses such as that they are robots or zombies.

Therefore, by inference to the best explanation, other people have minds.

The central problem with this inference is: What makes the hypothesis of other minds a better explanation than the alternatives?

According to the account of inference to the best explanation and explanatory coherence that I have developed elsewhere, the most important criteria for evaluating explanations are breadth (how much is explained), simplicity (how little is assumed), and analogy (how similar the explanations are to accepted ones) (see Thagard, 1988, 1992).

The hypothesis that there are other minds explains the full range of people's behavior without making a lot of special assumptions that would be required to say why people seem to have mental states when they are really robots or zombies. For example, the hypothesis that other people are robots remotely controlled by aliens requires assumptions about the existence and abilities of the controllers.

Moreover, the inference to other minds incorporates the analogical considerations mentioned earlier, because its overall plausibility is supported by the large number of similarities between my behaviors and those of others. The analogy is enriched by considering not only behavioral but also anatomical and physiological similarity. We know from neuroscience that most people's brains are very similar in both anatomy and functioning. If I wonder about whether you have a mind like mine and I make the reasonable inference that my mental states are closely tied to my brain states, then I can check out the similarities between you and me by having both of us undergo brain scans, with the predictable result that your brain is functioning like mine. Such findings strengthen further the analogy between myself and others, and contribute to the overall coherence and acceptability of my conclusion that you and other people have minds like mine. Just as I have brain states that produce my mental states and behaviors, so do you, by inference to the best explanation.

I think this solution to the skeptical problem of other minds is adequate, but it makes the problem harder than it needs to be if we understand how mirror neurons work. The best-explanation account requires us to have an abstract, verbal, conceptual representation of the mental states of others, opening up the possibility that this representation might be wrong. But mirror neurons provide a person with much more

direct understanding of what is going on in the mind of another person. When I see you pick up an apple, I do not need to construct a complex inference about what you are experiencing when you grasp it. Rather, my mirror neurons provide me with a mind-brain process that overlaps with the mind-brain process that occurs when I myself grasp something. Hence I do not have to construct an elaborate inference to understand your behavior: I just get it. Similarly, when I see you are pricked with a needle I do not need to infer by an analogy-based inference to the best explanation that you are feeling the same kind of pain that I would feel in a similar situation. Instead, my perception of your situation generates an internal brain process in cortical areas such as the insula and anterior cingulate which overlaps with the brain process that produces my own pain. I do not have to infer your pain – I actually feel something like your pain.

Mirror neurons are not by themselves a full solution to the problem of other minds, as there are many mental states, for example whatever you are thinking about this paper right now, that are not directly mirrored. But they do provide a useful supplement to the inference to the best explanation by providing some more physically direct connections that do not require complex inference. Similarly, although the philosophical problem of the existence of the external world is arguably solvable by an inference to the best explanation that there is a physical world that explains my sense experiences, we often do not have to make such inferences because of direct causal processes linking world and perception. In vision, for example, we know that light reflects off objects into our eyes, stimulating the neurons in the retina to send signals to layers of brain areas in the visual cortex. Our knowledge of the external world is based as much on this kind of direct causal process as it is on the verbal inference that there is a world independent of

the senses. Similarly, in mathematics our understanding of number is not a purely verbal construction, but rests on evolved neural processes for counting and comparing magnitudes that are also found in rats and pigeons (Dehaene, 1997). Just as mathematics grows out of our primitive number sense rooted in neurophysiology, and just as knowledge of the external world grows out of our perceptual ability to interact physically with the world, so our knowledge of other minds is partly underpinned by a kind of direct understanding via mirror neurons. I will return to the question of what this understanding amounts to at the end of this section.

The contention that mirror neurons help with the skeptical problem of other minds presupposes that the brain processes that enable people to mirror the mental states of others are generally reliable, producing accurate rather than erroneous conclusions about whether other people are experiencing pain, disgust, or other states. According to Goldman (forthcoming), mirror-based attribution of mental states to others is comparatively immune from error. Solution to the problem of other minds does not require infallibility in mental state ascription, and we do sometimes make mistakes in the inferences we make about the mental states of others. For example, people have a tendency to overestimate the extent to which the mental states of others are the same as their own. Such egocentric biases do not, however, intervene much in mental state attribution based on mirror systems. We can still be mistaken when mirroring leads us to attribute pain, disgust, or some other emotion to another person, who might be an actor pretending to be in such a state. However, such situations are rare, so it would appear that mirror-based assignment of mental states is generally reliable.

Just as mirror neurons contribute to a solution to the skeptical problem of other minds, they even more manifestly contribute to a solution to the social-cognition problem of how people manage, albeit imperfectly, to understand the minds of others, a process sometimes called *mentalizing* (Frith and Frith, 1999) or *mindreading* (Nichols and Stich, 2003). There are two main accounts of the nature of mentalizing, as theorizing and as simulating. On the theorizing account (sometimes called the theory theory), people understand others by having explicit causal accounts of why they behave as they do (Gopnik and Meltzoff, 1997). In contrast, on the simulation account, people often understand others by putting themselves in their place and running their own mental states to generate similar experiences and behaviors (Goldman, 2006).

I do not see these as competing theories, but rather as complementary accounts of different ways that people can understand each other (Barnes and Thagard, 1997). Under different circumstances, with different kinds of knowledge available, we will sometimes theorize and at other times simulate, and may combine these activities by using our theories to improve our simulations and vice versa. Neural mirroring is best viewed as a particularly direct non-verbal kind of simulation (Goldman, 2005; Gallese and Goldman, 1998).

Thus, contrary to the suggestion of Gallese, Keysers, and Rizzolatti (2004), mirror neurons do not provide a novel stand-alone alternative theory of mindreading, but rather a useful supplement to simulation accounts which are in turn a useful complement to the more easily grasped theory-based account of how we understand other people. Here is an integrated example. Suppose you see your friend John furiously pounding his fist on his desk and cursing. You can understand his action theoretically by applying some kind

of general rule or schema that people who are upset often behave in this way, or understand it via simulation by imagining yourself in John's situation (perhaps he just had a paper rejected) and simulating your physical reaction to the bad news, or both. More directly, if there are mirror neurons for anger as there are for disgust and pain, you will understand his situation by virtue of the fact that some of the neurons firing in your brain are the ones that fire when you are angry yourself. Hence the mirror-neuron account provides a useful supplement to theory-based and simulation-based accounts of how we understand other minds.

UNDERSTANDING

But does the activation of mirror neurons actually provide understanding? I think it does, but appreciating how requires an extension to familiar conceptions of explanation. Many philosophers and other theorists have maintained that explanations provide understanding of phenomena by locating them in causal networks (Salmon, 1984; Pearl 1988, 2000; Glymour, 2001; Thagard, 1999; Sloman, 2005). For example, we explain why someone got influenza by identifying all the causal factors – genetic makeup, environment, viral infection – that produced fever and the other symptoms of the disease.

In scientific discussions and in ordinary discourse, we normally use verbal descriptions to indicate causal relations, as in the theoretical account that John's paper rejection caused him to be angry which caused him to bang on his desk. I conjecture, however, that our fundamental understanding of causality is not verbal or mathematical. We know that A causes B is not just a matter of the probability of B given A being higher than the probability of B given not-A, for there may be other interacting factors

responsible for the increased probability. Drowning is more common on days when ice cream consumption is high, but that is not because ice cream causes drowning or drowning makes people want ice cream, but rather because hot weather leads to both more drowning and ice cream eating. Pearl (2000) and Woodward (2004) maintain that our understanding of causality is based on the idea of intervention, that A causes B is not just a probabilistic relation but depends on understanding that doing A produces B.

I conjecture that the idea of intervention is essentially sensorimotor rather than verbal. Every infant in a crib quickly learns that batting at a mobile or pushing a blanket will create a change in what is seen and touched: motor behavior produces visual and tactile changes. Perhaps there is an innate sensorimotor schema for this kind of causal intervention, or perhaps the human brain is just set up to acquire this schema quickly from early experience. Either way, the basic understanding of causality operates well before a child manages to acquire the verbal repertoire to talk about causes, let alone to reason about them using probability theory.

If the causal relation is fundamentally sensorimotor, then so must be representations of causes and their effects. The causal schema might be something like this:

<sensorimotor representation of action> -->

<sensorimotor representation of change>.

From the perspective of neuroscience, there is nothing mysterious about such nonverbal representations. Populations of neurons can encode aspects of the environment in many ways: visual, motor, other sensory, emotional, as well as verbal (see Eliasmith and Anderson, 2003, for a theory of neural representation). Now we can see how mirror

neurons provide understanding, by instantiating visual-motor representations of causal relations. When I observe you being pricked with a pin, it might activate my verbal pain schema: If an object pierces my skin, then I hurt. But in addition mirror neurons provide activation of a nonverbal schema in which the action and its result are both represented nonverbally, by populations of neurons that encode piercing and hurting through a combination of visual, motor, and other perceptual correlations. In this way I understand the causes of your behavior by a non-verbal representation of causality.

Identification of mental states that constitute crucial explanatory parts of the causal network can also be nonverbal. Just as my neural representation of my own pain is largely nonverbal, so my neural representation of someone else's pain derived from mirroring can also be nonverbal. Then the causal network that attributes a pain state to others would have a structure something like:

<sensorimotor representation of action> -->

<sensory representation of pain> -->

<sensorimotor representation of behavior>.

I apply this schema to myself in order to explain, for example, why hitting my thumb with a hammer causes me to scream and hold my thumb, but mirror neurons enable me to apply the same nonverbal schema to others.

Hence language and mirror neurons can both provide understanding by identifying a mental state and locating it in a causal network, but they do so in different ways because they use different kinds of representations of effects and causes and the relations between them. Given that monkeys and infants seem to have some comprehension of causality, I would argue that the nonverbal kinds of representations –

sensory, motor, etc. – are more fundamental to understanding causal relations than our verbal descriptions. It is in this sense that the understanding of other people via mirror neurons is more direct than our verbal versions of why people do what they do. This directness helps both with the skeptical problem of other minds and the social-cognition problem. In later sections, I will argue that it also helps with understanding empathy and moral motivation. At the end of this paper, I will flesh out this account of understanding by tying it to neurocomputational ideas about representation and inference.

EMPATHY

Empathy, where you imagine yourself in someone's situation and get some indication of their emotional state, is important for understanding other people and for making moral decisions about them. Barnes and Thagard (1997) presented an account of empathy as a kind of analogical mapping relying largely on verbal representations of someone's situation. Mirror neurons make possible a more direct kind of empathy employing visual-motor representations.

Barnes and Thagard (1997) schematize empathy in accord with the cognitive theory of emotions, according to which emotions are primarily indications of the extent to which personal goals are or are not being achieved (Oatley, 1992). For example, people are happy when their goals are being satisfied, and angry when someone or something is blocking the satisfaction of their goals. From this perspective, empathy has the following structure:

The person **P** is in situation **S**, which is like your situation **S'**.

P has goals **G** which are like your goals **G'**.

When you faced situation **S'** which affected your goals **G'**, you felt emotion **E'**. (**S'** and **G'** caused **E'**.)

So **P** is probably feeling emotion **E**, which is like your **E'**, caused by **S** and **G**.

For example, if you want to understand a friend who seems sad because of a disappointment, you can remember a situation such as a job rejection where you were sad because your career goals were not accomplished. Empathy is thus analogical mapping between someone else's situation and your own, and can be easily modeled computationally by ACME, a program which maps between similar situations and transfers information from one to the other (Holyoak and Thagard, 1989). A somewhat richer account of empathy as analogical mapping is accomplished by the model HOTCO, which incorporates ACME but allows the transfer of nonverbal representations of emotional valences corresponding to positive and negative emotions (Thagard, 2006). Still, HOTCO depends on highly verbal representations of the situations of both the provider and recipient of empathy.

In contrast, Singer et al. (2006) advocate a perception-action model of empathy, in which observation or imagination of another person in a particular emotional state automatically activates a representation of that state in the observer (see also Preston and de Waal, 2002; Decety and Jackson, 2004; Jackson et al., 2006; Singer et al., 2004). Singer and her colleagues used functional magnetic resonance imaging (fMRI) to measure brain activity in volunteers who observed others receiving painful stimulation to their hands. As expected, mere observation of another's pain produced increased activation in the pain network of the observer, including the insula and anterior cingulate.

As in the earlier study of Singer et al. (2004), people who scored higher on standard empathy scales had higher activity in these brain areas. It thus appears that more empathetic people have more active mirror neuron systems for appreciating the pain of others.

The major manipulation of the Singer et al. (2006) study was that the people who received painful stimulation had previously engaged in a game where some had behaved fairly and others unfairly. Men, and to a lesser extent women, showed much less pain-related brain activation for those sufferers who had acted unfairly. Moreover, men but not women showed greater activation in the reward-related area of the nucleus accumbens when observing unfair people being punished. Thus men more than women took pleasure in the pain of wrongdoers.

The studies of Singer and her colleagues (2004, 2006) suggest a need to expand the largely verbal account of empathetic analogical mapping provided by Barnes and Thagard (1997). The source analog is my own experience of what I experienced as the result of stimulation:

<visual/tactile representation of my stimulation> -->

<sensory/affective representation of my pain>

Then, when I see you stimulated similarly, the result is:

<visual representation of your stimulation> -->

<sensory/affective representation of my AND your pain>.

This mental operation is still a sort of analogical inference, in that it involves grasping a relational similarity between two situations. But it is much more direct than the verbal sort performed by ACME. The arrows indicate causality, but in a way that should not be

understood verbally. Rather, as suggested in the last section, causality itself is understood via sensorimotor neural representations. Thus feeling your pain can sometimes be a fairly direct reaction based on observation, not an intellectual exercise performed by seeing systematic mappings between two people's situations and goals. The intellectual, verbal kind of empathy may still occur, but it probably depends on the visual/motor/sensory/affective neural pathways that generate the emotional response that is the hallmark of empathy. I may feel your pain as the result of thinking about your situation and seeing parallels with my own experiences, but observing you in pain much more immediately gives me a sense of your pain.

Empathy as verbal analogy fits well with the cognitive theories of emotion advocated by Oatley (1992), Nussbaum (2001), and others; whereas empathy as physical experiences fits better with the physiological theories of emotion advocated by James (1884) and Damasio (1994). But a full theory of emotion should incorporate both cognitive and physiological processes (Thagard and Aubie, forthcoming), and so should an account of the full range of empathy. Hodges and Wegner (1997) review two complementary forms of empathy, one automatic and largely unconscious and the other controlled by conscious inferences.

My mirror-neuron account makes it clear how even the physically direct kind of empathy differs from emotional contagion, which involves picking up an emotion from someone else without any inference. According to Hatfield, Cacioppo, and Rapson (1994, pp. 10-11) theory of emotional contagion, one person "catches" another's emotions as follows:

1. In conversation, people tend automatically and continuously to mimic and synchronize their movements with the facial expressions, voices, postures, movements, and instrumental behavior of others.
2. Subjective emotional experiences are affected, moment to moment, by the activation and/or feedback from such mimicry.
3. Given propositions 1 and 2, people tend to “catch” others’ emotions, moment to moment.

In contrast, empathy via mirror neurons does not require mimicry or behavioral synchronization, but only the perception of another’s situation, which activates a kind of perceptual/motor schema that generates an analogous feeling. However, as de Vignemont and Singer (2006) out, even this kind of empathy can be modulated by appraisal processes.

In sum, empathy can be based on the kind of verbal analogical mapping discussed by Barnes and Thagard (1997), but more fundamentally can involve direct perceptual detection of the relation between someone’s situation and your own via your mirror neurons. Either way, the phrase “I feel your pain” is not just a touchy-feely cliché, but rather an expression of genuine appreciation of the experiences of others. Moreover, feeling the pain of others can contribute enormously to caring about them and being motivated to act ethically in general. Empathy is a major factor in the moral development of children (Hoffman, 2000).

MORAL MOTIVATION

Why be moral? This question is fundamental for ethics, because even if people can figure out what are the right things to do, we can ask why they would in fact do those

things. The problem of moral motivation – what makes people do what is right – has two classes of answers, rationalist and sentimentalist (Nichols, 2004). The traditional philosophical responses to the problem have been rationalist: We should be moral because it would be irrational to do otherwise. The rationality of morality might derive from a priori truths about what is right, as in Kantian deliberation, or from contractarian arguments that it is rational for people to agree with others to be moral, Nichols (2004) argues that a major problem for rationalism is that there is a class of people, psychopaths, who have no impediments in rationality but nevertheless see nothing wrong in harming other people.

Nichols argues convincingly that what is wrong with psychopaths is not their rationality but their emotions. Blair, Mitchell, and Blair (2005) review substantial evidence that psychopathy, whose symptoms include antisocial behavior, lack of guilt, and poverty of emotions, is the result of impairments to emotional learning that derive from disrupted functioning of the amygdala, an area of the brain well known to be important for processing emotions such as fear. That psychopathy derives from emotional problems fits well with the metaethical position of sentimentalism, according to which moral judgment is grounded in affective response. This tradition goes back to eighteenth-century writers such as David Hume and Adam Smith, and continues today (e.g. Gibbard, 1990; Nichols, 2004; Prinz, 2006).

According to Nichols (2004, p. 98), an adequate sentimental account must explain how emotion plays a role in linking moral judgment to motivation, while also allowing a place for reason in moral judgment. His explanation is cultural and historical: “Norms are more likely to be preserved in the culture if the norms resonate with our affective

systems by prohibiting actions that are likely to elicit negative affect.” (Nichols, 2004, p. 140). Norms that prohibit harm to others are virtually ubiquitous across cultures because of this “affective resonance.” The adoption of norms makes it possible to reason about what is right and wrong, but these norms have an emotional underpinning that intrinsically provides a connection between morality and action: people are moral because of their emotional commitment to normative rules.

What is missing from Nichols’ otherwise plausible account is an explanation of *why* people have such a basic emotional reaction to harm to others. There is no mystery concerning why you do not want harm to yourself, because experiences such as pain and fear are intrinsically negative. Appreciating harm to others might be done by the sort of abstract analogical reasoning discussed by Barnes and Thagard (1997), but there is no guarantee that such reasoning will be motivating: I may understand that you experience pain and fear, but why should I care? What makes emotional moral learning work?

As my discussion of empathy indicated, mirror neurons provide a plausible link between personal experience and the experience of others. People not only observe the pain and disgust of others, they experience their own versions of that pain and disgust, as shown by the mirroring activity in cortical regions such as the insula and anterior cingulate. Normal children do not need to be taught moral rules as abstract theological principles “Thou shalt not kill!” or rational ones “Act only in ways that could become universal.” Normal children do not need to reason about why harm is bad for other people: they can actually feel that harm is bad. Thus mirror neurons provide motivation not to harm others by virtue of direct understanding of what it is for another to be harmed.

It would be elegant if there were evidence that psychopaths have deficiencies in the functioning of their mirror neurons, but the relevant experiments have not yet been done. But it is possible that psychopaths' deficits in emotional learning, attributed by Blair, Mitchell, and Blair (2005) to disrupted functioning of the amygdala, are partly due to mirror neuron malfunctioning. Children who are incapable for genetic or environmental reasons of feeling the pain of others will not be able to become motivated to follow rules that direct them not to harm other people. Blair, Mitchell, and Blair (2005, p. 128) discuss moral socialization in terms of aversive conditioning, as when caregivers punish children for their wrongdoings. They claim that the sadness, fearfulness, and distress of a victim act as a stimulus to instrumental learning not to produce harm. The involvement of mirror neurons shows why instrumental learning can be especially effective when people can fully appreciate what is negative about their behavior.

I have argued that mirror neural mechanisms contribute to solution of the philosophical problem of moral motivation by showing how biologically normal people naturally have at least some understanding and concern for harm to other people. Feeling the pain of others is not the whole story of moral motivation, for there are many cognitive and social additions in the form of rules and expectations that can be built on top of neural mirroring. The motivating reason to be moral is not that just that morality is rational, but rather that feeling the pain of others is biologically part of being human.

A NEUROCOMPUTATIONAL ACCOUNT

In order to provide a more detailed account of the contribution of mirror neurons to understanding, I will now sketch a neurocomputational explanation of how they serve

to represent and make inferences about the actions and experiences of others. My account is based on the theory of neural representation developed by Chris Eliasmith (2003, 2005; Eliasmith and Anderson, 2003; see also Thagard forthcoming). Philosophers tend to understand representation in terms of verbal structures such as propositions, rules and concepts, but neural representations are more biologically fundamental. I will try to outline how mirror neurons can represent actions and mental states, and how they can contribute to understanding by constituting parts of action schemas that link perception, mental experience, and motor behavior.

On Eliasmith's account, the basic unit of representation is a neural population, which is a group of interconnected neurons whose firing patterns correlate systematically with events in the world. For example, there are groups of neurons in the fusiform gyrus (part of the brain's temporal lobe) that fire in response to faces; damage to this area can lead to prosopagnosia, the inability to identify different faces (Schiltz et al., 2005). A neural population can represent a set of faces by virtue of statistical dependencies (causal correlations) between its firing patterns and different faces. For example, we can hypothesize that whenever you see a picture of George W. Bush, the neurons in your fusiform gyrus fire in a regular pattern. A firing patterns does not mean just that some neurons are on and others are off, nor even just that some neurons are firing rapidly and others are firing slowly. Neurons encode information by means of spiking patterns constituted by temporal sequences of firing and not firing.

For neural representation, the correlation between a firing pattern in a population of neurons and what it stands for must be causal, based on a history of interactions between a neural population and what it represents. For example, if you have seen many

pictures of George W. Bush, the firing pattern in your fusiform gyrus is the result of the formation of synaptic connections among the neurons in a neural population. These synaptic connections can be excitatory or inhibitory: the firing of one neuron can encourage or discourage the firing of another. Given the perceptual input from a picture of Tony Blair and connections between neurons in a neural population, your fusiform gyrus can then recognize him by generating a pattern of firing that is very similar to ones that were generated by previous experiences of seeing him. For such firing patterns to be representational, the relevant neural population must have been tuned to Blair or pictures of him by previous causal interactions.

To apply this theory of neural representation to mirror neurons, we need to begin with an account of how people's neurons represent their own actions and experiences, and then move on to how they can represent the actions and experiences of others. A monkey can have a neural population represent its own actions by virtue of a causal correlation between a firing pattern of the neurons and occurrence of the action. For example, when a monkey raises its arm, a neural population fires in a particular way that is causally correlated with the arm raising. The neural pattern occurs just when the arm raising occurs, and co-occurrence is the result of previous causal interactions between arm raising and neural activity. Thus the neural pattern represents the monkey raising its arm.

As we saw earlier, what is remarkable about mirror neurons is that they fire in response to the arm raising of other monkeys. Thus the same neural population is able to represent both the firing of a monkey's arm and the raising of arms by other monkeys that it sees. How do these neurons represent an event – the other monkey raising its arm

– with which they have not been previously associated? The answer has partly to do with the natural way that neural populations generalize. A neural population can represent both particular objects and classes of objects. It becomes tuned to an object by repeatedly being given perceptual input generated by that object, but it becomes tuned to a class of objects by being given perceptual input generated by multiple similar objects. Once tuned to a class of objects, a neural population will display the firing pattern correlated with those objects when it encounters a new object with similar perceptual properties. Thus the population can identify a new object and represent it even if the population and object have not previously been causally correlated. Similarly, I conjecture, there is enough perceptual similarity between a monkey raising its own arm and another monkey raising its arm that the same neural population fires.

This account of action representation is independent of the question of whether the neural connections that generate firing patterns are innate or learned. My hunch is that they are largely learned, perhaps by a sophisticated kind of unsupervised learning of the sort described by Hinton, Osindero, and The (forthcoming). But it is also possible that the causal correlation between the firing patterns in a neural populations and the monkey's own actions is not the result of experience, but rather of evolutionary history: the monkey has inherited genes for producing neural connections that generate firing patterns correlated with particular actions. More plausible is some kind of interaction between innate structure and experiential learning, as seems to be the case for face recognition. Newborns have an immediate interest in faces, but it takes infants months of perceptual experience before they get very good at identifying people.

In sum, the monkey has neurons that represent its own actions by virtue of causal correlation with them, and the same neurons represent the actions of other monkeys by a kind of natural generalization. This recognition is a precursor to understanding, which requires action schemas of the sort discussed below. First it is necessary to show how the account I gave of monkey mirror neurons for representing actions can be applied to human mirror neurons for mental states such as pain.

The brains' representation of pain is highly complex, involving multiple areas such as the parietal-insular cortex, the anterior cingulate, and the amygdala (Craig, 2003). It would therefore be a gross oversimplification to associate a particular kind of pain with a localized neural population. More accurately, we can say that your experience of pain involves interrelated firing patterns of a group of neural populations. Nevertheless, the same idea of neural representation applies: a group of neural populations represents something by virtue of having firing activity that is causally correlated with what it represents. For example, your pain in stubbing your toe is produced by a pattern of activation in multiple neural populations. This pattern applies to multiple episodes of you stubbing a toe, so it has an element of generality to it, enabling you to identify and predict the experiences associated with new episodes. Thus your representation of your own mental experience – the pain of stubbing your toe – is a highly distributed pattern of neural firings. This neural representation is largely independent of the verbal representation of your experience that is constituted by neural activity in other brain areas dedicated to linguistic activity.

As we saw earlier, humans seem to have neural areas that mirror the pain of others, and it should now be clear how the pattern of neural activity that represents my

pain can also represent your being in pain. Just as the monkey's neural representation of its own arm raising is applied to arm raising by another monkey, so your neural representation of stubbing your toe is activated when you see someone else stubbing a toe, by a kind of recognition of similar perceptual inputs, including the motion of the foot, the banging of the toe, and the behavioral response. Hence I recognize your pain in a way very similar to how I recognize or predict new instances of my own pain, not by verbal inference but by generation of a pattern of neural firing.

We now have a sketch of the basic mechanism underlying the operation of mirror neurons, but still lack an account of how they furnish a kind of non-verbal understanding. To provide this, we need a neural account of inference as well as representation. From a neural perspective, inference looks quite different from the picture familiar from formal logic, in which inference is derivation of a proposition from other propositions by means of rules such as modus ponens. In contrast, an inference in a neural system should be viewed as a transformation in patterns of neural firing. When a neural population is firing in a particular way, and its firing pattern changes because of interactions with perceptual inputs or other neural populations, then an inference has been made. This change can either be transitory, applying just to the current situation, or more permanent if synaptic changes occur that constitute learning and lead to differences in future patterns of neural firing. As in verbal inference systems such as logic, inference in neural systems can either be particular (inferring a new singular proposition) or general (inferring a new generalization).

I argued earlier that explanation consists of locating something puzzling in a causal network, and now it is possible to see how neural populations can represent such

networks. Representing that A causes B is more complex than just *if A then B*, or even of a transformation of the neural representation of A into a neural representation of B. Rizzolatti (2005b) points out that mirror neurons enable an organism to locate the experiences of other organisms into a larger pattern of perceptual inputs, perceptions, and behaviors. I will use the term *action schema* to refer to a structure of the form:

sensory inputs -> mental experience -> motor behavior.

The arrow here represents causality based on the kind of sensorimotor experience I discussed earlier. Everyone has numerous action schemas, for example based on experiences of seeing an object flying toward your head, feeling fear, and ducking. In neural terms, these schemas are represented by transformations in neural patterns, including ones for sensory inputs, mental experience, and motor action.

Stringing together multiple action schemas can enable someone to simulate more complex situations, involving multiple actions. For example, I can imagine you being hit on the head with a stick, grimacing in pain, reacting in anger, and then punching the person who hit you. My simulation of your behavior can use the kinds of verbal representations in the last sentence, but it can also use nonverbal action schemas involving neural representations of visual, tactile, auditory, emotional, and pain experiences. I agree with Goldman (2006, forthcoming) that mirror neurons are not the only way of simulating the mental processes of others, and that simulation is not the only way of attributing mental states. But it is important to appreciate that at least some of the time mirror neurons enable us to tap into neural processes involving sensorimotor representations in order to appreciate the behavior and internal states of others without verbal theorizing.

When I understand my own experiences, I do so by locating them in action schemas that may be verbal and explicit but may equally well be non-verbal at the neural level of representations of sensation, experience, and action. For example, I understand the pain of stubbing my toe by locating it in the causal nexus that includes my perceptions of walking through a room, the experience of the pain, and my actions in jumping back from the obstacle that my toe hit and rubbing the toe.

Mirror neurons make it possible for a similar kind of understanding to apply to others. When I see you stub your toe, I can explain what is happening verbally by rules such as “If people hit an obstacle, they can stub their toe” and “If people stub their toe, it hurts.” However, in addition to verbal rules that may constitute an explicit action schema, people also have implicit action schemas that operate in neural populations. Thus my perception of you stubbing your toe produces the firing pattern in my mirror neurons for pain which is similar to the firing pattern they evince when I have stubbed my own toe. My brain can then apply the same input-experience-behavior action schema that I apply in my own case to your case as well. This is how mirror neurons provide a more direct kind of understanding than more familiar verbal explanations.

Such understanding is still inferential, but not in the familiar sense in which analogy and inference to the best explanation are well known kinds of inference. These require verbal representations and very high-level kinds of reasoning. On the neural approach I have been advocating, inference is understood not just as manipulations of verbal symbols, but more generally as transformation of patterns of neural firing. The relevant patterns can constitute sensorimotor representations as well as verbal symbols. Verbal explanations are less direct in that they require the translation of sensorimotor

inputs into verbal symbols, which then can be used to construct causal networks. The process looks something like this:

(1) sensory inputs -> activation of verbal representation of causal network

-> attribution of mental state represented verbally

In contrast, the more direct but still inferential process that mirror neurons make possible looks something like:

(2) sensory inputs -> activation of nonverbal representations of causal network

-> implicit attribution of mental state represented nonverbally

The second process is still inferential in the sense of requiring transformation of neural patterns, so it would be a mistake to call it immediate or automatic. But dropping out the need to translate sensorimotor representations into verbal ones makes it possible to locate observations in a causal network without verbal elaboration. A deeper understanding of how this works will require the development, now in the planning stage, of new computational models of biologically realistic neural networks that can accomplish such representations and transformations.

Hence when I infer that you have the mental state of pain, I am justified for two different kinds of reasons. The more familiar philosophical reason is that the best explanation of your behavior after stubbing your toe is that you are in pain. But my mirror neurons, responding to my perception of you stubbing your toe, serve to represent your pain as well as mine, as part of my learned action schema. Hence I am justified in believing that you are in pain independently of the philosophical inference, which adds an extra layer of justification to the more commonplace one furnished by mirror neurons. Moreover, the understanding furnished by mirror neurons has the value that it can

generate empathy, enabling a person to feel something like what someone else is feeling, thereby increasing moral motivation.

CONCLUSION

Hence the functioning of mirror neurons contributes to explanation of other minds, empathy, and moral motivation. This functioning does not by itself constitute a solution to any of these problems, but shows how part of the solution in each case is a more direct causal connection that complements the more traditional solutions based on conscious, verbal inferences. Such complementation does not make sense if one views justification as akin to verbal deduction, but fits well with a view of justification as coherence involving parallel satisfaction of multiple constraints, including ones represented non-verbally (Thagard, 2000). Coherentist justification of this sort is the result of multiple parallel processes all converging on a common conclusion, showing how nonverbal mental state attribution based on mirror systems can harmonize with verbal mental state attributions based on analogy and inference to the best explanation.

Mirror neurons are neither necessary nor sufficient for understanding the actions and mental states of others. We can have understanding without the operation of mirror neurons, because we are often able to locate the behavior and mental behavior of others in causal networks that are represented verbally and visually in parts of the brain not involved in mirroring. Mirror neurons by themselves are not sufficient to understand the actions and experiences of others because they handle only part of the neural representation of action schemas, which also include sensory inputs and motor behaviors. Nevertheless, mirror neurons can make an important contribution to understanding others by representing a crucial part of action schemas in non-verbal explanations.

The current discussion is only a preliminary to future work in both neuroscience and philosophy. For neuroscience, it would be desirable to have a much more detailed account of how mirroring works, that is how some groups of neurons respond both to one's own experience and to one's perception of others. Perhaps it will be possible to expand current models of cognitive/affective integration to indicate how one person's emotional state can also be used to appreciate the emotional state of another (Wagar and Thagard, 2004; Litt, Eliasmith, and Thagard, 2006; Thagard and Aubie, forthcoming). Such improved models should be able to give a much deeper understanding of the philosophically important problems of other minds, empathy, and moral motivation.

Acknowledgements. For comments on an earlier draft, I am grateful to Lorraine Besser-Jones, Chris Eliasmith, Abninder Litt, and several anonymous referees. This research was supported by the Natural Sciences and Engineering Research Council of Canada.

REFERENCES

- Avenanti, A., Bueti, D., Galati, G., & Aglioti, S. M. (2005). Transcranial magnetic stimulation highlights the sensorimotor side of empathy for pain. *Nature Neuroscience*, 8(7), 955-960.
- Barnes, A., & Thagard, P. (1997). Empathy and analogy. *Dialogue: Canadian Philosophical Review*, 36, 705-720.
- Blair, J., Mitchell, D. R., & Blair, K. (2005). *The psychopath: Emotion and the brain*. Malden, MA: Blackwell.
- Botvinick, M., Jha, A. P., Bylsma, L. M., Fabian, S. A., Solomon, P. E., & Prkachin, K. M. (2005). Viewing facial expressions of pain engages cortical areas involved in the direct experience of pain. *Neuroimage*, 25(1), 312-319.
- Craig, A. D. (2003). Pain mechanisms: Labeled lines versus convergence in central processing. *Annual Review of Neuroscience*, 26, 1-30.
- Damasio, A. R. (1994). *Descartes' error*. New York: G. P. Putnam's Sons.
- de Vignemont, F., & Signer, T. (2006). The empathic brain: how, when and why? *Trends in Cognitive Science*, 10, 435-440.
- Decety, J., & Jackson, P. L. (2004). The functional architecture of human empathy. *Behavioral and Cognitive Neuroscience Reviews*, 3(2), 71-100.
- Dehaene, S. (1997). *The number sense*. New York: Oxford University Press.
- di Pellegrino, G., Fadiga, L., Fogassi, L., Gallese, V., & Rizzolatti, G. (1992). Understanding motor events: a neurophysiological study. *Experimental Brain Research*, 91(1), 176-180.

- Eliasmith, C. (2003). Moving beyond metaphors: Understanding the mind for what it is. *Journal of Philosophy*, 100, 493-520.
- Eliasmith, C. (2005). Neurosemantics and categories. In H. Cohen & C. Lefebvre (Eds.), *Handbook of categorization in cognitive science* (pp. 1035-1054). Amsterdam: Elsevier.
- Eliasmith, C., & Anderson, C. H. (2003). *Neural engineering: Computation, representation and dynamics in neurobiological systems*. Cambridge, MA: MIT Press.
- Frith, C. D., & Frith, U. (1999). Interacting minds: A biological basis. *Science*, 286, 1692-1695.
- Gallese, V., Fadiga, L., Fogassi, L., & Rizzolatti, G. (1996). Action recognition in the premotor cortex. *Brain*, 119 (Pt 2), 593-609.
- Gallese, V., & Goldman, A. I. (1998). Mirror neurons and the simulation theory of mindreading. *Trends in Cognitive Sciences*, 2, 493-501.
- Gallese, V., Keysers, C., & Rizzolatti, G. (2004). A unifying view of the basis of social cognition. *Trends in Cognitive Sciences*, 8(9), 396-403.
- Gibbard, A. (1990). *Wise choices, apt feelings*. Cambridge, MA: Harvard University Press.
- Glymour, C. (2001). *The mind's arrows: Bayes nets and graphical causal models in psychology*. Cambridge, MA: MIT Press.
- Goldman, A. I. (2005). *Mirror systems, social understanding and social cognition*. Retrieved May 17, 2006, from <http://interdisciplines.org/mirror/papers/3>
- Goldman, A. I. (2006). *Simulating minds*. New York: Oxford University Press.

- Goldman, A. I. (forthcoming). Mirroring, mindreading, and simulation. In J. A. Pineda (Ed.), *Mirror neuron systems: The role of mirroring processes in social cognition*. Totowa, NJ: Human Press.
- Gopnik, A., & Meltzoff, A. N. (1997). *Words, thoughts, and theories*. Cambridge, MA: MIT Press.
- Graham, G. (1998). *Philosophy of mind: An introduction* (2nd ed.). Oxford: Blackwell.
- Hatfield, E., Cacioppo, J. T., & Rapson, R. L. (1994). *Emotional contagion*. Cambridge: Cambridge University Press.
- Hinton, G. E., Osindero, S., & Teh, Y. (forthcoming). A fast learning algorithm for deep belief nets. *Neural Computation*.
- Hodges, S. D., & Wegner, D. M. (1997). Automatic and controlled empathy. In W. Ickes (Ed.), *Empathic accuracy* (pp. 311-339). New York: The Guilford Press.
- Hoffman, M. L. (2000). *Empathy and moral development: Implications for caring and justice*. Cambridge: Cambridge University Press.
- Holyoak, K. J., & Thagard, P. (1989). Analogical mapping by constraint satisfaction. *Cognitive Science*, *13*, 295-355.
- Hyslop, A. (2005). *Other minds*. Retrieved May 17, 2006, from <http://plato.stanford.edu/archives/win2005/entries/other-minds>
- Iacoboni, M., Molnar-Szakacs, I., Gallese, V., Buccino, G., Mazziotta, J. C., & Rizzolatti, G. (2005). Grasping the intentions of others with one's own mirror neuron system. *PLoS Biology*, *3*(3), e79.

- Jackson, P. L., Brunet, E., Meltzoff, A. N., & Decety, J. (2006). Empathy examined through the neural mechanisms involved in imagining how I feel versus how you feel pain. *Neuropsychologia*, *44*(5), 752-761.
- James, W. (1884). What is an emotion? *Mind*, *9*, 188-205.
- Kohler, E., Keysers, C., Umiltà, M. A., Fogassi, L., Gallese, V., & Rizzolatti, G. (2002). Hearing sounds, understanding actions: action representation in mirror neurons. *Science*, *297*(5582), 846-848.
- Litt, A., Eliasmith, C., & Thagard, P. (2006). Why losses loom larger than gains: Modeling neural mechanisms of cognitive-affective interaction. In R. Sun & N. Miyake (Eds.), *Proceedings of the twenty-eighth annual meeting of the Cognitive Science Society* (pp. 495-500). Mahwah, NJ: Erlbaum.
- Nelissen, K., Luppino, G., Vanduffel, W., Rizzolatti, G., & Orban, G. A. (2005). Observing others: multiple action representation in the frontal lobe. *Science*, *310*(5746), 332-336.
- Nichols, S. (2004). *Sentimental rules : on the natural foundations of moral judgment*. Oxford: Oxford University Press.
- Nichols, S., & Stich, S. (2003). *Mindreading: An integrated account of pretense, self-awareness, and understanding other minds*. Oxford: Oxford University Press.
- Nussbaum, M. (2001). *Upheavals of thought*. Cambridge: Cambridge University Press.
- Oatley, K. (1992). *Best laid schemes: The psychology of emotions*. Cambridge: Cambridge University Press.
- Pearl, J. (1988). *Probabilistic reasoning in intelligent systems*. San Mateo: Morgan Kaufman.

- Pearl, J. (2000). *Causality: Models, reasoning, and inference*. Cambridge: Cambridge University Press.
- Preston, S. D., & de Waal, F. B. (2002). Empathy: Its ultimate and proximate bases. *Behav Brain Sci*, 25(1), 1-20; discussion 20-71.
- Prinz, J. (2006). The emotional basis of moral judgments. *Philosophical Explorations*, 9, 29-43.
- Rizzolatti, G. (2005a). The mirror neuron system and imitation. In S. Hurley & N. Chater (Eds.), *Perspectives on imitation: From neuroscience to social science. Volume 1: Mechanisms of imitation and imitation in animals* (pp. 55-76). Cambridge, MA: MIT Press.
- Rizzolatti, G. (2005b). The mirror neuron system and its function in humans. *Anatomy and Embryology*, 210, 419-421.
- Rizzolatti, G., & Arbib, M. A. (1999). From grasping to speech: imitation might provide a missing link: reply. *Trends in Neurosciences*, 22(4), 152.
- Rizzolatti, G., & Craighero, L. (2004). The mirror-neuron system. *Annual Review of Neuroscience*, 27, 169-192.
- Rizzolatti, G., Fadiga, L., Gallese, V., & Fogassi, L. (1996). Premotor cortex and the recognition of motor actions. *Cognitive Brain Research*, 3(2), 131-141.
- Salmon, W. (1984). *Scientific explanation and the causal structure of the world*. Princeton: Princeton University Press.
- Schiltz, C., Sorger, B., Caldara, R., Ahmed, F., Mayer, E., Goebel, R., et al. (2006). Impaired face discrimination in acquired prosopagnosia is associated with

- abnormal response to individual faces in the right middle fusiform gyrus.
Cerebral Cortex, 16(4), 574-586.
- Singer, T., Seymour, B., O'Doherty, J., Kaube, H., Dolan, R. J., & Frith, C. D. (2004).
Empathy for pain involves the affective but not sensory components of pain.
Science, 303(5661), 1157-1162.
- Singer, T., Seymour, B., O'Doherty, J. P., Stephan, K. E., Dolan, R. J., & Frith, C. D.
(2006). Empathic neural responses are modulated by the perceived fairness of
others. *Nature*, 439(7075), 466-469.
- Sloman, S. A. (2005). *Causal models: How people think about the world and its
alternatives*. New York: Oxford University Press.
- Thagard, P. (1988). *Computational philosophy of science*. Cambridge, MA: MIT
Press/Bradford Books.
- Thagard, P. (1992). *Conceptual revolutions*. Princeton: Princeton University Press.
- Thagard, P. (1999). *How scientists explain disease*. Princeton: Princeton University
Press.
- Thagard, P. (2000). *Coherence in thought and action*. Cambridge, MA: MIT Press.
- Thagard, P. (2006). *Hot thought: Mechanisms and applications of emotional cognition*.
Cambridge, MA: MIT Press.
- Thagard, P. (forthcoming). How cognition meets emotion: Beliefs, desires, and feelings
as neural activity. In G. Brun, U. Doguoglu & D. Kuenzle (Eds.), *Epistemology
and emotions*. Aldershot: Ashgate.

Thagard, P., & Aubie, B. (forthcoming). Emotional consciousness: A neural model of how cognitive appraisal and somatic perception interact to produce qualitative experience. *Consciousness and Cognition*.

Wagar, B. M., & Thagard, P. (2004). Spiking Phineas Gage: A neurocomputational theory of cognitive-affective integration in decision making. *Psychological Review*, *111*, 67-79.

Wicker, B., Keysers, C., Plailly, J., Royet, J. P., Gallese, V., & Rizzolatti, G. (2003). Both of us disgusted in my insula: The common neural basis of seeing and feeling disgust. *Neuron*, *40*(3), 655-664.

Woodward, J. (2004). *Making things happen: A theory of causal explanation*. Oxford: Oxford University Press.