REVIEW

Predictive coding: an account of the mirror neuron system

James M. Kilner · Karl J. Friston · Chris D. Frith

Abstract Is it possible to understand the intentions of other people by simply observing their actions? Many believe that this ability is made possible by the brain's mirror neuron system through its direct link between action and observation. However, precisely how intentions can be inferred through action observation has provoked much debate. Here we suggest that the function of the mirror system can be understood within a predictive coding framework that appeals to the statistical approach known as empirical Bayes. Within this scheme the most likely cause of an observed action can be inferred by minimizing the prediction error at all levels of the cortical hierarchy that are engaged during action observation. This account identifies a precise role for the mirror system in our ability to infer intentions from actions and provides the outline of the underlying computational mechanisms.

Keywords Mirror neurons · Action observation · Bayesian inference · Predictive coding

Introduction

The notion that actions are intrinsically linked to perception was proposed by William James, who claimed, "every mental representation of a movement awakens to some degree the actual movement which is its object" (James 1890). The implication is that observing, imagining, or in anyway representing an action excites the motor program

J. M. Kilner (✉) · K. J. Friston · C. D. Frith
The Wellcome Trust Centre for Neuroimaging,
Institute of Neurology, 12 Queen Square,
WC1N 3BG London, UK
e-mail: j.kilner@fil.ion.ucl.ac.uk

used to execute that same action (Jeannerod 1994; Prinz 1997). Interest in this idea has grown recently, in part due to the neurophysiological discovery of "mirror" neurons. Mirror neurons discharge not only during action execution but also during action observation, which has led many to suggest that these neurons are the substrate for action understanding.

Mirror-neurons were first discovered in the premotor area, F5, of the macaque monkey (Di Pellegrino et al. 1992; Gallese et al. 1996; Rizzolatti et al. 2001; Umilta et al. 2001) and have been identified subsequently in an area of inferior parietal lobule, area PF (Gallese et al. 2002; Fogassi et al. 2005). Neurons in the superior temporal sulcus (STS), also respond selectively to biological movements, both in monkeys (Oram and Perrett 1994) and in humans (Frith and Frith 1999; Allison et al. 2000; Grossman et al. 2000) but they are not mirror-neurons, as they do not discharge during action execution. Nevertheless, they are often considered part of the mirror neuron system (MNS; Keysers and Perrett 2004) and we will consider them as such here. These three cortical areas, which constitute the MNS, the STS, area PF and area F5, are reciprocally connected. In the macaque monkey, area F5 in the premotor cortex is reciprocally connected to area PF (Luppino et al. 1999) creating a premotor–parietal MNS and STS is reciprocally connected to area PF of the inferior parietal cortex (Harries and Perrett 1991; Seltzer and Pandya 1994) providing a sensory input to the MNS (see Keysers and Perrett 2004 for a review). Furthermore, these reciprocal connections show regional specificity. Although STS has extensive connections with the inferior parietal lobule, area PF is connected to an area of the STS that is specifically activated by observation of complex body movements. An analogous pattern of connectivity between premotor areas and inferior parietal lobule has also been

demonstrated in humans, both directly (Rushworth et al. 2006) and indirectly (Iacoboni et al. 2001, 2005). In addition, a sequential pattern of activation in the human MNS has been demonstrated during action–observation that is consistent with the proposed pattern of anatomical connectivity (Nishtani and Hari 2000, 2002).

Mirror-neurons and the MNS have been the focus of much interest since their discovery because they have been proposed as a neural substrate that could enable us to understand the intentions of others through the observation of their actions (Gallese and Goldman 1998). Actions can be understood at many different levels. After Hamilton and Grafton (2007) here we will consider actions that can be described at four levels. (1) The intention level that defines the long-term goal of an action. (2) The goal level that describes short-term goals that are necessary to achieve the long-term intention. (3) The kinematic level that describes the shape of the hand and the movement of the arm in space and time. (4) The muscle level that describes the pattern of muscle activity required to execute the action. Therefore to understand the intentions or goals of an observed action, the observer must be able to describe the observed movement at either the goal level or the intention level having only access to a visual representation of the kinematic level. Although mirror neurons have been proposed as the neural substrate that could enable us to understand the intentions or goals of an observed action (Gallese and Goldman 1998) little is known about the neural mechanisms underlying this ability to 'mind read'. Gallese (2006) recently noted that '' ... we do not have a clear neuroscientific model of how humans can understand the intentions promoting the actions of others they observe''. Therefore, the question remains, if mirror-neurons do mediate understanding of actions done by others how do they do it?

Rizzolatti and Craighero (2004) suggested, ''The proposed mechanism is rather simple. Each time an individual sees an action done by another individual, neurons that represent that action are activated in the observer's premotor cortex. This automatically induced, motor representation of the observed action corresponds to that which is spontaneously generated during active action and whose outcome is known to the acting individual. Thus, the mirror-neuron system transforms visual information into knowledge''.

Generative and recognition models

Although this proposed mechanism is 'simple' in conception it is non-trivial in terms of implementation. It is unclear how the visual information from an observed action maps onto the observer's own motor system and how the goal of that action is inferred (Gallese et al. 2004; Iacoboni

2005; Jacob and Jeannerod 2005; Saxe 2005). Implicit in this and many descriptions of the MNS is the idea that visual information is transformed as it is passed by forward connections along the MNS network from low-level representations of the movement kinematics to high-level representations of intentions subtending the action. In this scheme, the observation of an action drives the firing of neurons in the STS, which drives activity in area PF, which in turn drives activity in area F5 (Fig. 1a). Formally, this is a recognition model that operates by the inversion of a generative model, where the generative model produces a sensory representation of the kinematic level of an action given the information at the goals or intentions level. Generative models can be framed in terms of a deterministic non-linear generative function
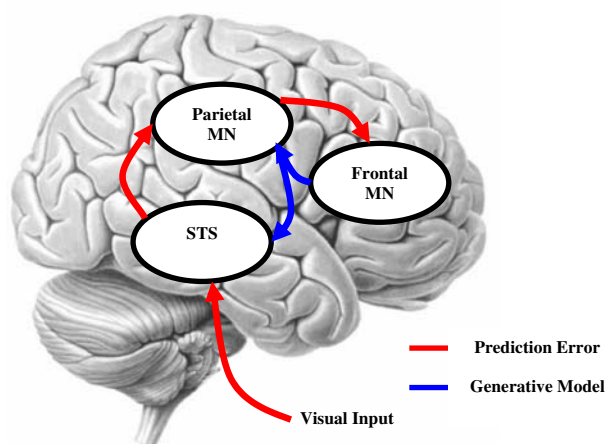
$$u = G(v, \theta).$$

Here v is a vector (i.e., a list) of underlying causes and u represents some sensory inputs. $G(v, \theta)$ is a function that

a) The MNS as a feedforward recognition model
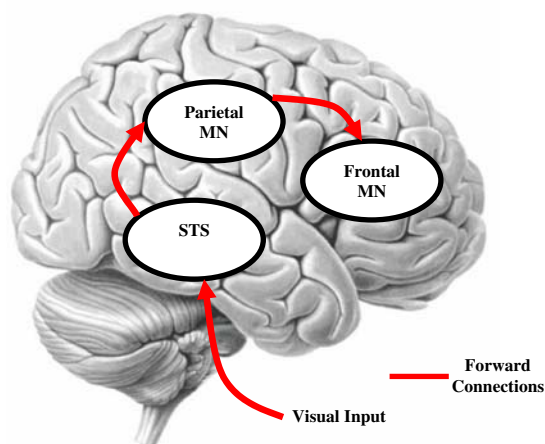
b) The MNS as a predictive coding model

Fig. 1 Schemas of the mirror-neuron system

generates inputs from the causes given some parameters of the model, θ. In the case of action execution/observation the causes, v are the long-term intentions or goals of the action. The parameters, θ, correspond to the connection strengths in the brain's model of how the inputs are caused. These are fixed quantities that have to be learned, presumably through development. The inputs, u is the visual signal corresponding to the sight of the executed action. This generative model will produce an estimate of the visual consequence of an executed action given the cause or goals of that action. By inverting this generative model it is possible to infer the cause or goals of an action given the visual input. However, is there any evidence that such generative models exist?

All executed actions have a sensory consequence. For example when we reach and grasp a bottle of wine there will be a change in our proprioceptive signal as we move, there will be a change in our tactile signal as we touch the bottle of wine, there will be change in the visual signal as we observe the action we are executing and there may be a change in the auditory signal as we pick up the bottle of wine. It is now generally accepted that when we execute a movement we predict the sensory consequences of that movement through generative or forward models (Wolpert et al. 1995, 2003; Wolpert and Miall 1996). These predictions can then be used to finesse motor control problems induced by delayed feedback and sensory noise. In short, forward models that generate predicted kinematics from motor commands are considered an integral part of motor execution. The suggestion here is that these generative models can be inverted to infer the causes given the data.

One of the obvious problems with such a model is that this scheme will only work when the processes generating the sensory inputs from the causes are invertible, i.e., when one sensory input is associated uniquely with one cause. In general, this is not the case since the same sensory input can have many causes. In the specific case of action–observation the same kinematics can be caused by different goals and intentions. For example, if, while walking along the street, someone suddenly waves his arm, is he hailing a taxi or swatting a wasp? A trivial example of this is given by the generative model $u = v^2$. In this example knowing u does not uniquely determine v, which could be negative or positive. The nature of this ill-posed problem has been demonstrated empirically in the MNS. Mirror-neurons in area F5 that discharge when a monkey is observing a reach and grasp action also discharge when the sight of end point of this movement is occluded (Umilta et al. 2001). Critically, this result shows that mirror-neurons in area F5 are not simply driven by the visual representation of an observed movement. Therefore, if the inversion of a generative model is not sufficient to explain how we can understand others' actions through observation, then how can this be achieved?
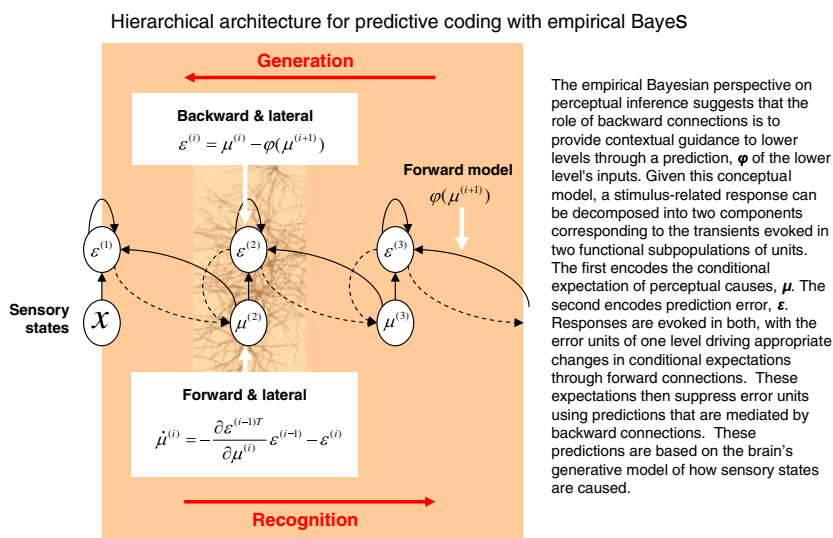
The question remains, if mirror-neurons do mediate understanding of actions done by others how do they do it?

Predictive coding and the MNS

The perspective we propose here is that the role of the mirror-neuron system in reading or recognising the goals of observed actions can be understood within a predictive coding framework. Predictive coding is based on minimizing prediction error though recurrent or reciprocal interactions among levels of a cortical hierarchy (Box 1). In the predictive coding framework, each level of a hierarchy employs a generative model to predict representations in the level below. This generative model uses backward connections to convey the prediction to the lower level where it is compared to the representation in this subordinate level to produce a prediction error. This prediction error is then sent back to the higher level, via forward connections, to adjust the neuronal representation of sensory causes, which in turn change the prediction. This self-organising, reciprocal exchange of signals continues until prediction error is minimised and the most likely cause of the input has been generated. It can be shown that this scheme is formally equivalent to empirical Bayesian inference, in which prior expectations emerge naturally from the hierarchal models employed (see Box 2; Friston 2002, 2003, 2005). It should be noted that the prediction addressed in predictive coding is predicting the sensory effects from their cause. This is about the mapping between the cause (motor commands to grasp) and the sensory (i.e., visual or proprioceptive) expression or effect of that cause. It is not about forecasting (i.e., predicting the sensory states in the future, given the sensory state now), aka prospective coding (see Schultz-Bosbach and Wolfgang Prinz 2007 this issue for a review of this topic).

For the MNS this means that anatomically the areas engaged by movement observation are arranged hierarchically and the anatomical connections between these areas are reciprocal. In terms of functional anatomy it means that the prediction error encoding higher-level attributes will be expressed as evoked responses in higher cortical levels of the MNS. For action observation the essence of this approach is that, given a prior expectation about the goal of the person we are observing, we can predict their motor commands. Given their motor commands we can predict the kinematics on the basis of our own action system. The comparison of this predicted kinematics with the observed kinematics generates a prediction error. This prediction error is used to update our representation of the person's motor commands (Fig. 1b). Similarly, the inferred goals are updated by minimising the prediction error between the predicted and inferred motor

Box 1

Hierarchical architecture for predictive coding with empirical Bayes

Generation

Backward & lateral
$$\varepsilon^{(i)} = \mu^{(i)} - \varphi(\mu^{(i+1)})$$

Forward model
$$\varphi(\mu^{(i+1)})$$

Sensory states

$\varepsilon^{(1)}$ $\varepsilon^{(2)}$ $\varepsilon^{(3)}$

x $\mu^{(2)}$ $\mu^{(3)}$

Forward & lateral
$$\dot{\mu}^{(i)} = -\frac{\partial \varepsilon^{(i-1)T}}{\partial \mu^{(i)}} \varepsilon^{(i-1)} - \varepsilon^{(i)}$$

Recognition

The empirical Bayesian perspective on perceptual inference suggests that the role of backward connections is to provide contextual guidance to lower levels through a prediction, φ of the lower level's inputs. Given this conceptual model, a stimulus-related response can be decomposed into two components corresponding to the transients evoked in two functional subpopulations of units. The first encodes the conditional expectation of perceptual causes, μ. The second encodes prediction error, ε. Responses are evoked in both, with the error units of one level driving appropriate changes in conditional expectations through forward connections. These expectations then suppress error units using predictions that are mediated by backward connections. These predictions are based on the brain's generative model of how sensory states are caused.

Predictive coding and empirical Bayes

Predictive coding is a framework that, in a hierarchical setting, is equivalent to empirical Bayesian inference. Schemes for identifying the causes of sensory input that are based entirely on bottom-up, forward connections, such as the feedforward recognition model in Figure 1a, are ill-posed when the generative model linking sensations and causes can not be inverted. This can occur when there are inherent ambiguities in the way sensations are generated (*e.g.*, visual occlusion) However, identification is possible for noninvertible generative models by incorporating constraints and prior information. This is equivalent to full Bayesian inference. Although full Bayes enables recognition with noninvertible generative models it creates a new problem for the brain. Namely, the brain cannot construct prior expectations *de novo*; these have to be learned and adapted to the current experiential context. Statistical solutions are available for this problem using empirical Bayesian inference, in which priors are estimated from data. Empirical Bayesian inference harnesses the hierarchical structure of a generative model, treating the estimates at one level as priors on the subordinate level. This provides a natural framework within which to treat cortical hierarchies in the brain, each level providing constraints on the level below. This approach models the world as a hierarchy of systems where supraordinate causes induce and moderate changes in subordinate causes. These priors offer contextual guidance towards the most likely cause of the sensory input. This scheme can be implemented with simple architectures that have a degree of biological plausibility (see Box 3)

Box 2

commands (see Box 1). By minimizing the prediction error at all the levels of the MNS, the most likely cause of the action will be inferred at all levels (intention, goal, motor and kinematic). This approach provides a mechanistic account of how responses in the visual and motor systems are organised and explains how the cause of an action, can be inferred from its observation.

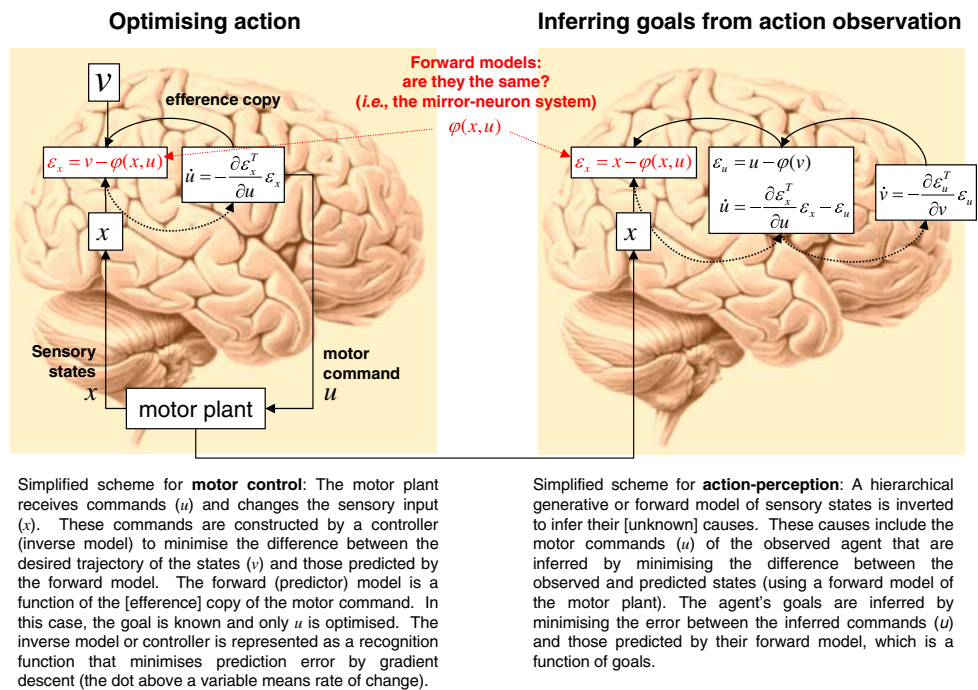## Generative models in motor control

Predictive coding is particularly appropriate for understanding the function of the MNS; predictive coding provides an established computational framework for inferring the causes (intentions, goals and motor commands) of sensory inputs (observed kinematics). It is now generally accepted that forward or generative models play a critical role in motor control (Wolpert et al. 1995, 2003; Wolpert and Miall 1996). The suggestion here is that these same models are used to infer motor commands from observed kinematics produced by others during perceptual inference (see Chater and Manning 2006) for a similar heuristic in the domain of language perception). Box 3 illustrates the formal similarities and differences between action–optimisation and action–perception. In execution, motor commands are optimised to minimise the difference between predicted and desired kinematics, under the assumption that the desired kinematics (i.e., goals) are known. Conversely, in action–perception, these goals have to be inferred. However, in both optimisations a forward model of motor control is required. In the predictive coding account of the MNS, the same generative model used to predict the sensorial effects of our own actions can also be used (with appropriate transformations) to predict the actions of others (see Friston 2005 for a description of the relationship between forward and inverse models and predictive coding).

There have been several previous accounts that have proposed the use of forward and inverse models in action–observation (Keysers and Perrett 2004; Wolpert et al. 2003; Miall 2003). ''Skilled motor behaviour relies on the brain learning both to control the body and predict the consequences of this control. Prediction turns motor commands into expected sensory consequences, whereas control turns desired consequences into motor commands. To capture this symmetry, the neural processes underlying prediction and control are termed the forward and inverse internal models, respectively'' (Flanagan et al. 2003). First, forward

Box 3

Optimising action

Inferring goals from action observation

Forward models:
are they the same?
(*i.e.*, the mirror-neuron system)
$\varphi(x,u)$

v

efference copy

$\varepsilon_x = v - \varphi(x,u)$ 　 $\dot{u} = -\dfrac{\partial \varepsilon_x^T}{\partial u}\varepsilon_x$

x

Sensory
states
x

motor
command
u

motor plant

$\varepsilon_x = x - \varphi(x,u)$ 　 $\varepsilon_u = u - \varphi(v)$ 　 $\dot{v} = -\dfrac{\partial \varepsilon_u^T}{\partial v}\varepsilon_u$

$\dot{u} = -\dfrac{\partial \varepsilon_x^T}{\partial u}\varepsilon_x - \varepsilon_u$

x

Simplified scheme for **motor control**: The motor plant receives commands (u) and changes the sensory input (x). These commands are constructed by a controller (inverse model) to minimise the difference between the desired trajectory of the states (v) and those predicted by the forward model. The forward (predictor) model is a function of the [efference] copy of the motor command. In this case, the goal is known and only u is optimised. The inverse model or controller is represented as a recognition function that minimises prediction error by gradient descent (the dot above a variable means rate of change).

Simplified scheme for **action-perception**: A hierarchical generative or forward model of sensory states is inverted to infer their [unknown] causes. These causes include the motor commands (u) of the observed agent that are inferred by minimising the difference between the observed and predicted states (using a forward model of the motor plant). The agent's goals are inferred by minimising the error between the inferred commands (u) and those predicted by their forward model, which is a function of goals.

and inverse models have been proposed as an account of imitation; the inverse model (mapping kinematics to motor signals) is identical to feedforward recognition model of the MNS (shown in Fig. 1a). The logic is that the inverse model can used as recognition model and therefore infer the cause of an observed action. Once the cause of the observed kinematics is inferred the action can then be imitated. Second, the HMOSIAC model for motor control has recently been proposed as a model for understanding social interactions (Wolpert et al. 2003). The links between this model and the predictive coding account exist at a number of levels. In the HMOSIAC model several of predictor–controller pairs are organised hierarchically. The predictor (forward) model is employed to predict the input in a subordinate module and the controller is used to adjust the predictor to maximise the prediction.

Although these generalisations of forward–inverse models in motor control to imitation and social interactions are exciting, they are formally distinct from, and more complicated than, the predictive coding account of the MNS. In predictive coding there is no separate inverse model or controller; a forward model is simply inverted by suppressing the prediction error generated by the forward model. This inversion depends on the self-organising, reciprocal exchange of signals between hierarchical levels (see Box 1). This simplicity translates into an algorithmic architecture that could be implemented plausibly by the brain (and for which there is a considerable amount of anatomical and physiological evidence). Indeed Miall (2003) when describing the HMOSIAC model wrote,

''Quite how this multi-level controller could be generated neurally is not yet clear, but the link with mirror-neurons seems appealing''. In contrast, the predictive coding account has been described in some detail at the neural level (see Friston 2002, 2003, 2005).

An example of the predictive coding account of the MNS

Within predictive coding, recognition of causes is simply the process of jointly minimizing prediction error at all levels of a cortical hierarchy. The most likely cause of an observed action (i.e., motor commands, goal or intention) can be estimated from the visual representation of the observed movement. An intuitive example is given in Fig. 2. Here we use the predictive coding account of the MNS to address the Dr. Jekyll and Mr Hyde thought-experiment described in Jacob and Jeannerod (2005). In this thought-experiment, one is invited to watch identical movements, made by Dr. Jekyll and Mr Hyde. In both cases one observes the same person, taking hold of a scalpel and applying it to a human body. However, in one case Dr. Jekyll is using the scalpel to cure a patient but in the other Mr Hyde's aim is to inflict pain. Jacob and Jeannerod argue that the MNS is incapable of distinguishing between these two intentions, as the observed movement is identical in both cases. This is certainly true for the bottom–up inverse or recognition model described in Fig. 1a but it is not true for the predictive coding scheme. The observed kinematics
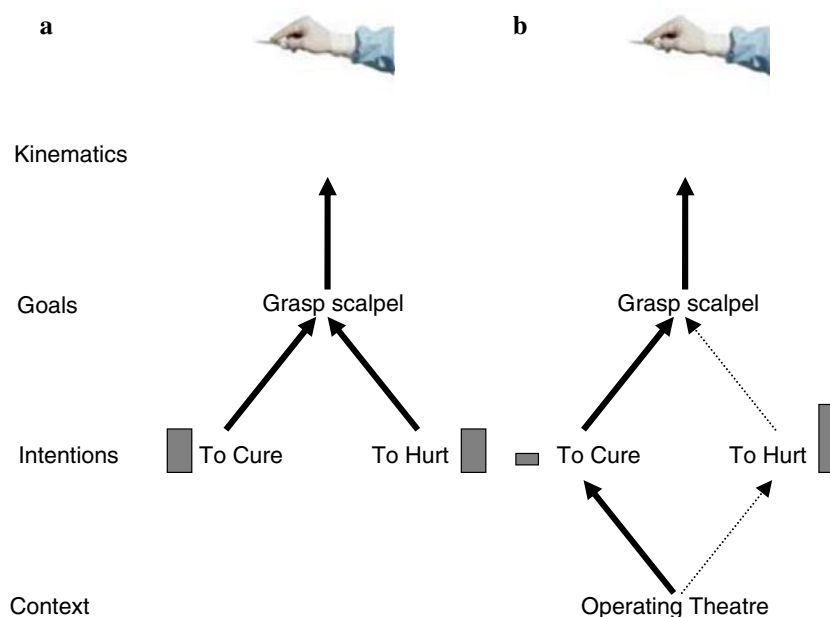
Fig. 2 Examples of the predictive coding account of the MNS. Here we consider four levels of attribution in an example hierarchy of the MNS; kinematics, goal, intention and context. In **a** action–observation is considered in the absence of a context, in **b** the identical action is observed in but now in the context of an operating theatre. The *bars* depict the level degree of prediction error. In **a** both intentions predict identical goals and kinematics and therefore the prediction error is identical in both schemes. In this case the model can not differentiate between the intentions causing the action. In **b** the context causes a large prediction error for the goal 'to hurt' and a small prediction error for the goal 'to cure'. In this case the model can differentiate between the two intentions

can be explained at a number of levels that are hierarchically organised, the visual representation of the kinematics, the underlying motor signals, the short-term goal (e.g., to grasp the scalpel), and the long-term intention ('to cure' or 'to hurt'). These three levels are shown schematically in Fig. 2a. In predictive coding, the intentional level predicts a goal that in turn predicts the kinematic representation of the motor acts. At each level the predicted activity is compared to the actual activity and any difference is projected back-up the hierarchy as a prediction error (see Box 1). In the case where both intentions produce identical movements there are identical prediction errors and therefore the predictive coding account can not infer a unique intention from the observed movement. However, in contradistinction to the bottom–up model, the predictive coding model also has to explain sensory information pertaining to the context in which the movement has been observed. This induces high-level sensory causes that provide empirical priors on action–perception; for example, a therapeutic intention explains the action and the visual scenery, if seen in an operating theatre (Fig. 2b). This does not mean that context is coded by mirror-neurons but rather the MNS is part of a larger hierarchy, where intentions are encoded. In this scheme, the intention that is inferred from the observation of the action now depends upon the prior information received from a context level. In other words, if the action was observed taking place in an

operating theatre there would be a large prediction error for the intention 'to hurt' and a smaller prediction error for the intention 'to cure'. The prediction error would be the same at all other levels of the hierarchy for the two intentions. By minimising the overall prediction error the MNS would infer that the intention of the observed movement was to cure. Therefore, the MNS is capable of inferring a unique intention even if two intentions result in identical movements. This observation is supported empirically. Mirror-neurons in area PF have been shown to have differential patterns of firing when viewing movements that are virtually identical at the kinematic level, but differ at the level of intention. In this task there is a contextual cue, the object that is grasped, that informs the monkey of the intention of the action to be observed. Within the predictive coding account the MNS will always be able to infer the most likely intention of an observed action, given the observer's priors.

Summary

Social interaction depends upon our ability to infer beliefs and intentions in others. Impairments of this ability can lead to major developmental and psychiatric disorders such as autism (Dapretto et al. 2006; Oberman et al. 2005) and schizophrenia (Arbib and Munhenk 2005). It has been

suggested that the MNS could underlie this ability to 'read' someone else's intentions. Here we have proposed that the MNS is best considered within a predictive coding framework. One of the attractions of predictive coding is that it can explain how the MNS could infer someone else's intentions through observation of their movements. Within this scheme the most likely cause of an observed action is inferred by minimising the prediction error at all levels of the cortical hierarchy that is engaged during action–observation. Central to testing the predictive coding account of the MNS is that the nodes of the cortical hierarchy are well characterised both anatomically and functionally. From the existing literature we can assume that any MNS network will include areas of ventral premotor cortex, inferior parietal lobule and STS. However, the function of each node in the MNS and the hierarchical organisation of the MNS are not known. Implicit in many accounts of the MNS is the notion that the area F5 is the highest level of the hierarchy. This is the hierarchical arrangement shown in Fig. 1. However, there is no direct evidence to support this view and the results of recent studies suggest that the inferior parietal lobule area may be superordinate to premotor areas in the MNS hierarchy (Hamilton and Grafton 2006; Fogassi et al. 2005). Specifically, the theory underlying the predictive coding account of the MNS is independent of the hierarchical organisation. The predictive coding account of the MNS specifies a precise role for the MNS in our ability to infer intentions and formalises the underlying computations. It also connects generative models that are inverted during perceptual inference with forward models that finesse motor control.

References

Allison T, Puce A, McCarthy G (2000) Social perception from visual cues: role of the STS region. Trends Cogn Sci 4:267–278

Arbib MA, Mundhenk TN (2005) Schizophrenia and the mirror-neuron system: an essay. Neuropsychologia 43:268–280

Chater N, Manning CD (2006) Probabilistic models of language processing and acquisition. Trends Cogn Sci 10(7):335–344

Dapretto M, Davies MS, Pfeifer JH, Scott AA, Sigman M, Bookheimer SY, Iacoboni M (2006) Understanding emotions in others: mirror-neuron dysfunction in children with autism spectrum disorders. Nat Neurosci 9:28–30

Di Pellegrino G, Fadiga L, Fogassi L, Gallese V, Rizzolatti G (1992) Understanding motor events: a neurophysiological study. Exp Brain Res 91:176–180

Flanagan JR, Vetter P, Johansson RS, Wolpert DM (2003) Prediction precedes control in motor learning. Curr Biol 13(2):146–150

Fogassi L, Ferrari PF, Gesierich B, Rozzi S, Chersi F, Rizzolatti G (2005) Parietal lobe: from action organization to intention understanding. Science 308:662–667

Friston KJ (2002) Funtional integration and inference in the brain. Prog Neurobiol 68:113–143

Friston KJ (2003) Learning and inference in the brain. Neural Netw 16:1325–1352

Friston K (2005) A theory of cortical responses. Philos Trans R Soc Lond B Biol Sci 360:815–836

Frith CD, Frith U (1999) Interacting minds—a biological basis. Science 286:1692–1695

Gallese V (2006) Embodied simulation: from mirror neuron systems to interpersonal relations. In: Bock G, Goode J (eds) Empathy and fairness. Novartis Foundation, San Diego

Gallese V, Goldman A (1998) Mirror-neurons and the simulation theory of mind reading. Trends Cogn Sci 2:493–501

Gallese V, Fadiga L, Fogassi L, Rizzolatti G (1996) Action recognition in the premotor cortex. Brain 119:593–609

Gallese V, Fogassi L, Fadiga L, Rizzolatti G (2002) Action representation and the inferior parietal lobule. In: Prinz W, Hommel B (eds) Attention and performance XIX. Common mechanisms in perception and action. Oxford University Press, Oxford, pp 247–266

Gallese V, Keysers C, Rizzolatti G (2004) A unifying view of the basis of social cognition. Trends Cogn Sci 8:396–403

Grossman E, Donnelly M, Price R, Pickens D, Morgan V, Neighbor G, Blake R (2000) Brain areas involved in perception of biological motion. J Cogn Neurosci 12:711–720

Hamilton AF, Grafton ST (2006) Goal representation in human anterior intraparietal sulcus. J Neurosci 26:1133–1137

Hamilton AF, Grafton ST (2007) The motor hierarchy: from kinematics to goals and intentions. In: Rosetti Y, Kawato M, Haggard P (eds) Attention and performance xxii (in press)

Harries MH, Perrett DI (1991) Visual processing of faces in temporal cortex: physiological evidence for a modular organization and possible anatomical correlates. J Cogn Neurosci 3:9–24

Iacoboni M (2005) Neural mechanisms of imitation. Curr Opin Neurobiol 15:632–637

Iacoboni M, Koski LM, Brass M, Bekkering H, Woods RP, Dubeau MC, Mazziotta JC, Rizzolatti G (2001) Reafferent copies of imitated actions in the right superior temporal cortex. Proc Natl Acad Sci USA 98(24):13995–13999

Iacoboni M, Molnar-Szakacs I, Gallese V, Buccino G, Mazziotta JC, Rizzolatti G (2005) Grasping the intentions of others with one's own mirror neuron system. PLoS Biol 3(3):e79

Jacob P, Jeannerod M (2005) The motor theory of social cognition: a critique. Trends Cogn Sci 9:21–25

James W (1890) Principles of psychology. Holt, New York

Jeannerod M (1994) The representing brain–neural correlates of motor intention and imagery. Behav Br Sci 17:187–202

Keysers C, Perrett DI (2004) Demystifying social cognition: a Hebbian perspective. Trends Cogn Sci 8:501–507

Luppino G, Murata A, Govoni P, Matelli M (1999) Largely segregated parietofrontal connections linking rostral intraparietal cortex (areas AIP and VIP) and the ventral premotor cortex (areas F5 and F4). Exp Brain Res 128:181–187

Miall RC (2003) Connecting mirror neurons and forward models. Neuroreport 14(17):2135–2137

Nishitani N, Hari R (2000) Temporal dynamics of cortical representation for action. Proc Natl Acad Sci USA 97(2):913–918

Nishitani N, Hari R (2002) Viewing lip forms: cortical dynamics. Neuron 36(6):1211–1220

Oberman LM, Hubbard EM, McCleery JP, Altschuler EL, Ramachandran VS, Pineda JA (2005) EEG evidence for mirror-neuron dysfunction in autism spectrum disorders. Brain Res Cogn Brain Res 24:190–198

Oram MW, Perrett DI (1994) Responses of anterior superior temporal polysensory (STPa) neurons to biological motion stimuli. J Cogn Neurosci 6:99–116

Prinz W (1997) Perception and action planning. Eur J Cogn Psych 9:129–154

Rizzolatti G, Craighero L (2004) The mirror-neuron system. Annu Rev Neurosci 27:169–192

Rizzolatti G, Fogassi L, Gallese V (2001) Neurophysiological mechanisms underlying the understanding and imitation of action. Nat Rev Neurosci 2:661–670

Rushworth MFS, Behrens TEJ, Johansen-Berg H (2006) Connection patterns distinguish three regions of human parietal cortex. Cereb Cortex 16:1418–1430

Saxe R (2005) Against simulation: the argument from error. Trends Cogn Sci 9:174–179

Schütz-Bosbach and Wolfgang Prinz (2007) Prospective coding in event representation (this issue)

Seltzer B, Pandya DN (1994) Parietal, temporal, and occipital projections to cortex of the superior temporal sulcus in the rhesus-monkey: a retrograde tracer study. J Comp Neurol 343:445–463

Umilta MA, Kohler E, Gallesse V, Fogassi L, Fadiga L, Keysers C, Rizzolatti G (2001) I know what you are doing. A neurophsyiological study. Neuron 31:155–165

Wolpert DM, Doya K, Kawato M (2003) A unifying computational framework for motor control and social interaction. Philos Trans R Soc Lond B Biol Sci 29:358:593–602

Wolpert DM, Ghahramani Z, Jordan MI (1995) An internal model for sensoriymotor integration. Science 269:1880–1882

Wolpert DM, Miall RC (1996) Forward models for physiological motor control. Neural Netw 9:1265–1279