Indeed, all functions of goal-states are necessary elements of the processes described by the SCM, although they are not yet properly analyzed. When current input is confronted with a target-state, the system is assumed to know what features of the environment are adequate as targets. But how does it know? In mechanical artifacts, for example, the boiler-thermostat system, the target is set by another agent. In adaptive biological systems, some basic target-states are "set" by evolution, but most of the goals relevant for the system are actually self-determined. Hence, some questions arise: How does the agent select a state of affairs as an appropriate target? Why do we use certain results as frame of reference for action control, rather than others?

The answers require invoking the evaluative function of goals: Starting from simple biologically determined goals (e.g., nutrition, proximity to conspecifics), the agent starts appraising positively certain environmental configurations (e.g., food on the table), while perceiving others to be negative (e.g., being isolated). The former are selected as potential target-states, whereas the latter are labeled as things to be avoided, possibly generating other target-states to ensure that they are indeed avoided. These more specific target-states may then acquire evaluative autonomy, if they prove successful over time in satisfying the agent's basic needs.

Here it is worth emphasizing that evaluation and motivation are *not* exclusive features of personal-level goal-states (desires and intentions). On the contrary, these functions apply also to basic, subpersonal needs, as discussed. Hence, appeals to the functional level of explanation favored by Hurley cannot get the SCM off the hook of this criticism. Evaluation and motivation characterize both personal and subpersonal goal-states, and yet they go unnoticed in the SCM.

Significantly, the target-state is the only component that never changes its dynamics through the five layers of the SCM. This is tantamount to saying that the model does not care for goal generation and goal revision. It being so, can the SCM really claim to enable deliberation and goal-based action understanding? Can we be satisfied with the characterization of deliberation as comparison of alternative predictions – Millikan's "trials and errors in the head" (Millikan 2006) – while nothing is said about how we compare different goals to choose our future conduct? The SCM reduces deliberation to deciding *how* to achieve a given end, whereas it remains silent on deciding *what* to achieve. This is a badly maimed picture of human deliberation: whenever decisions are made, the *what* is at least as important as the *how*.

As for understanding the intentional structure of action, this implies appreciating that intentions justify teleologically the observed behavior (i.e., motivating and controlling it; Dennett 1987), and that success or failure in achieving goals matters to the agent, that is, the world will be affectively appraised against them (Frijda 1986; Ortony et al. 1988). Moreover, mindreading should provide information on the goal of the observed action, that is, the target-state in the SCM (Gallese & Goldman 1998; Gallese et al. 1996). But there is no evidence that the model is capable of explaining anything of the sort: Mirroring, paired with action inhibition at layer 4, associates input with covert motor activation, which in turn associates with simulative prediction of the next input. Nothing of this, however, associates with the target-state, that is, the proper goal the system should recognize in the action of another agent. So it is doubtful that the SCM, in its present form, describes the informational resources enabling goal-based action understanding.

How could the SCM be amended to rectify these shortcomings? First, it should accommodate the possibility that *multiple targets* are active, so that facing the same input might yield divergent pressures on the agent's conduct. Even simple organisms confront similar dilemmas, when a certain action (e.g., eating in the open) can have conflicting results (e.g., foraging vs. exposure to predators). Without allowing for multiple targets, the SCM could never claim to enable proper deliberation, that is, choosing the ends as well as the means. Second, the *target-state must be incorporated in the dynamic loop of action and perception* to provide a grasp on goal dynamics (Castelfranchi & Paglieri 2007). Targets are abandoned, either in favor of better options or because they are satisfied (cyclically or permanently), while other targets emerge, either for instrumental reasons or because something unexpected and rewarding interests the agent. This requires supplementing the comparator mechanism with new functions, including the possibility of registering a mismatch between actual input and target *without* taking action – because evaluation of the world need not always trigger an attempt to change it. Third, the SCM must modify its interpretation of mirroring to account for goal understanding: To this purpose, *target-states must be adequately included in the mirror circuit*.

Ironically, the moral seems to be that the embodied/enactive perspective on social cognition should take the problem of "mind detachment" more seriously. If the aim is to prove that higher cognitive processes are grounded in bodily actions, better arguments are needed to show how mental states, goals among them, become increasingly *detached from actions*, in phylogenesis as much as in ontogenesis (Pezzulo & Castelfranchi 2007; Tomasello 1999). This is indeed a noble quest – one that the SCM has successfully pioneered, but still falls short of having completed.

ACKNOWLEDGMENTS

We are grateful to Maria Miceli, Giovanni Pezzulo, and Luca Tummolini for insightful discussion.

Putting the subjective back into intersubjective: The importance of person-specific, distributed, neural representations in perception-action mechanisms

doi: 10.1017/S0140525X07003287

Stephanie D. Preston

Department of Psychology, University of Michigan, Ann Arbor, MI 48109.

prestos@umich.edu http://www.umich.edu/~prestos

Abstract: The shared circuits model (SCM) relies on well-regarded theories of perception-action, mirror neurons, and forward models, but the functional/informational level of the model limits its ability to explain complex behavior such as true imitation. Data from our lab and others confirm the more general details of the model, accepted by most, but specify the neural mechanisms involved in perception-action processes.

The shared circuits model (SCM) has much in common with existing models of complex behavior and relies on some known properties of the nervous system. For example, most researchers no longer hold a pure version of the "sandwich model" and assume that perception and action overlap at the level of representation. Similarly, most agree that imitation exists on a continuum, with complex forms of true imitation relying on and evolving from more simple forms of reflexive imitation. In addition, there is general agreement that mirror neurons and forward models are relevant to questions about how we bridge the intersubjective divide and model others as we do ourselves. Therefore, most of Hurley's theoretical review is consistent with existing theory and data on the mechanisms of complex interpersonal phenomena, making it unlikely that anyone will take issue with the basic premises of the model; on the flip side, this also means that the model is limited in its ability to stimulate new directions for the field.

The five-layer model is the unique contribution of the model. However, I think that this part of the model suffers from being pitched at the functional/informational level. This is unfortunate and unnecessary given that much of the theory relies on very specific mechanisms for motor control and perspective taking that are precisely defined and empirically supported. Hence, the model could have been aimed at a functional neuroanatomical level, which would have made it more specific and more accurate.

The model also seems both underspecified and overspecified. In places where the literature is most agnostic on how certain processes work, Hurley is also agnostic. For example, the model intermixes concrete and abstract concepts (such as "targets," which can be either motor goals or life plans) without specifying whether we use the same neural processes for both, or just use analogous processes when planning to reach for a cup or to overthrow the government. In contrast, the formulation of the model into five discrete layers seems ill-fated, limiting the ability of the model to accord with the structure and functions of the nervous system. For example, layers 1 and 2 likely overlap a great deal in the brain because both require the cerebellum and act in concert to control action (cf. Wolpert et al. 1998). Conversely, there is no reason from phylogeny or ontogeny to assume that these two layers of control are primary to or evolved before the mirroring mechanisms of layer 3. Layer 2 focuses on visual and tactile feedback from the periphery, which are actually slow forms of feedback that forward models were designed to surpass. Layer 4 focuses almost entirely on "monitored inhibition" to segregate activation related to self and other, but it is unclear which type of inhibition is inculcated here (spinal, brain stem, frontal?), and there are many other ways in which self and other activation can be differentiated. Thus, it seems that there are ambiguities and inconsistencies in the model that could have been rectified by making more specific reference to the existing data on how the brain processes information.

Our lab seeks to understand the ways in which people process and understand the emotions of others. Like Hurley, we believe that basic emotion processing and related intersubjective phenomena, such as empathy, rely on an evolutionary conserved and basic perception-action mechanism (PAM) whereby perception of the emotional state of another automatically activates one's own representations for the state and situation (Preston & de Waal 2002).

Supporting Hurley's general rejection of the sandwich model, functional imaging work on empathy has found overlap between self and other processes in regions associated with subjective feeling states (Jackson et al. 2006; Lamm et al. 2007; Preston et al. 2007; in preparation b; Singer et al. 2004; 2006). Further supporting Hurley's application of perception-action processes to simulation, we have also found almost complete overlap in the neural substrates associated with imagining a personal past emotional experience and "trying on" the experience of another subject; however, we also found differences in self and other processes, which would not be predicted by the SCM, but are explicit in the PAM (Preston et al. 2007). In this study, the overall level of brain activation and autonomic arousal were much higher in the self-condition than the other-condition, and subjects recruited additional regions of visual association cortex when imagining another's scenario. These data suggest that online simulation of actual, personal events can differ in both quality and quantity from that of hypothetical events. Importantly, however, we found these differences between self and other only when subjects *could not* relate well to the situation of the other; there were no differences in neural patterns or autonomic arousal when subjects selected scenarios to which they could relate strongly (Preston et al. 2007).

This latter interaction reflects an important and overlooked point about the processing of other's actions and states: Perception-action mechanisms *require* that the subject have an existing representation for the action or state of the other. Thus, monkeys do not have mirror neurons for hand manipulations they do not understand, babies do not imitate gestures that they cannot make, and people cannot resonate with an unfamiliar emotional state and cannot predict your response to a truly novel situation.

We have striking pilot data to support this emphasis on personal representations, as individuals with depression perceive and respond to the distress of others differently than their non-depressed counterparts – they are less personally distressed by the sadness and hopelessness of hospital patients, and they are more likely to feel empathy and offer help to patients with particularly high need (Preston et al., in preparation a).

In another behavioral study, we have found that the mere perception of an emotional facial expression not only activates mirroring in a subject's facial muscles (cf. Dimberg & Oehman 1996), and primes the same valence in the subject (cf. Murphy & Zajonc 1993), but also rapidly activates the *semantic-level* representation for the *specific* emotion (e.g., "fear") (Preston & Stansfield, in press) – this finding is not predicted by models that exclusively rely on motor-based, facial feedback, or mirroring of emotion processing, but it is obvious from basic facts about how information is processed from perception to concept retrieval.

It is exciting and promising to have many researchers agreeing on some basic tenants about how behavior is instantiated – however, as with all complex problems, the devil is in the details. In order to make additional headway from here on out, we must look to the data.

ACKNOWLEDGMENTS
I am indebted to Frans de Waal for encouraging this piece and to R. Brent Stansfield for his feedback and support while writing it.

In search of a conceptual location to share cognition

doi: 10.1017/S0140525X07003299

Gün R. Semin[a] and John T. Cacioppo[b]
[a]*Faculty of Social Sciences, Utrecht University, 3584 CS Utrecht, The Netherlands;* [b]*Department of Psychology, The University of Chicago, Chicago, IL 60637.*
GR.Semin@psy.vu.nl **http://www.cratylus.org**
Cacioppo@uchicago.edu
http://psychology.uchicago.edu/people/faculty/cacioppo/

Abstract: It is argued that the multilayered model offered by the shared circuits model (SCM) falls short of capturing an essential aspect of social cognition, namely, its distributed nature. The SCM therefore falls short of modeling emergent social cognition and behavior.

Disciplinary perspectives cut the same realities in different ways, and so it is with philosophy, psychology, neuroscience, inter alia. At times, these disciplined languages cross their linguistic barriers and reach out to systematize knowledge at a level that supersedes the specific limitations imposed by their indigenous language and competence. And, even then, the particular slice of reality that we focus upon is conditioned by presuppositions about the nature of the beast we are examining.

The shared circuits model (SCM) provides a tour de force of portraying a multilayered account of *social cognition*, which is somewhat specifically grounded on imitation, whereby imitative learning is seen as a sophisticated form of social cognition. While *social* cognition appears to be a central construct for the SCM, the entire model is focused on an analysis at the *individual* level. It is this aspect of the SCM that we intend to complement in our commentary by drawing attention to the importance of distributed processes taking place between two or more individuals and the emergent quality of *social* cognition. Indeed, much of