# The Concept of Simulation in Control-Theoretic Accounts of Motor Control and Action Perception

**Mitchell Herschbach (mherschb@ucsd.edu)**
Department of Philosophy, University of California, San Diego, 9500 Gilman Drive #0119
La Jolla, CA 92092-0119 USA

## Abstract

Control theory is a popular theoretical framework for explaining cognitive domains such as motor control and "mindreading." Such accounts frequently characterize their "internal models" as "simulating" things outside the brain. But in what sense are these "simulations"? Do they involve the kind of "replication" simulation (R-simulation) found in the simulation theory of mindreading (Goldman, 2006)? I will argue that some but not all control-theoretic appeals to "simulation" involve R-simulation. To do so, I examine in detail a recent computational model of motor control and action perception based in control theory (Oztop et al., 2005). I argue that the architecture does not use R-simulation during motor control, but does during action perception. A novel result of this analysis is that the forward model—the control-theoretic mechanism most often described as performing simulation—is not well characterized in terms of R-simulation. I conclude with some lessons for research on the mechanisms of mindreading.

**Keywords:** Action Perception; Control Theory; Mindreading; Motor Control; Simulation Theory; Theory Theory

## Introduction

The term "simulation" is frequently used in cognitive science research, but often without specifying its meaning. One context where "simulation" has been fairly well defined is the *simulation theory* (ST) of mindreading (i.e., our ability to form beliefs or representations about people's mental states). In Alvin Goldman's (2006) recent defense of ST, he specifies two meanings of "simulation": (1) *replication* simulation (R-simulation), the kind described by ST; and (2) a *computational modeling* sense of simulation (CM-simulation), akin to the theorizing characteristic of the rival *theory theory* (TT) of mindreading.

"Simulation" also comes up in applications of *control theory* (CT) to cognitive domains such as motor control and mindreading—in particular, action perception (i.e., understanding the goal or intention of a perceived action). These control architectures use "internal models" described as "simulating" things outside the brain. But it is not always clear what is meant here by "simulation." As Goldman (2006) shows, some descriptions suggest that CT-architectures involve R-simulation and are, accordingly, applications of ST; but others suggest the "simulation" involved is CM-simulation, and thus not consistent with ST. Goldman's assessment is that CT's "simulations" are better characterized in terms of CM-simulation than R-simulation. Yet Goldman admits that his conclusion is tentative: that a synthesis of CT and ST remains an open possibility, but that "the issue awaits full resolution" (p. 217).

In this paper I defend a limited synthesis of ST and CT: that some but not all applications of CT involve R-simulation. To do so I will discuss in detail one CT-based computational model of motor control and action perception (Oztop et al., 2005). The model characterizes a mechanism for visually-guided motor control, which is modified for the purpose of action perception. I argue that this architecture's use for motor control is best characterized in terms of CM-simulation, but that its use for action perception does, contrary to Goldman's initial assessment, involve R-simulation. A surprising result of this analysis is that the forward model—the control theoretic mechanism commonly described as performing simulation—is not actually the part best characterized as performing R-simulation. I end with some implications of this analysis for research into the mechanisms of mindreading.

## Two Senses of "Simulation"

Goldman's (2006) two senses of simulation, R- and CM-simulation, each concern "models" which "simulate" some target system, typically for the purpose of explanation and prediction. They differ with regard to the nature of the relation between a simulation and what it represents.

CM-simulation, as its name suggests, has its home in the computational models constructed in order to better understand real-world systems. Such models are described as "simulating" processes of a target system when "the computer generates correct symbolic descriptions of its outputs from descriptions of its inputs by means of descriptions of its intervening states" (p. 35). CM-simulations represent the processes of a target system, but do not "work according to the same principles, or undergo the same (or even similar) states, as the simulated system" (p. 35). One way a CM-simulation could work is to use a *theory* describing the target system, and make inferences about the evolution of the system by applying the theory. For example, a CM-simulation of a weather system may work by applying theoretical generalizations describing weather systems (presumably, laws of physics) in order to represent the causal processes by which the weather system evolves over time. But this simulation process is not itself directly governed by the same causal processes as, and does not enter identical/similar states as, the weather system. If it is run on a computer, for instance, the simulation is instead driven by the syntactic rules constituting the program. Whatever its physical implementation (in a brain, computer, etc.), a CM-simulation represents the inputs, outputs, and intervening states of a target system. The simulation process

only has this semantic relation to the target system's processes—it does not resemble or duplicate the target process in other ways (e.g., physically).

R-simulations, however, serve their representational role in just this way. An R-simulation is a process with the function of replicating or resembling a target system. For example, a scale physical model of our solar system can replicate or resemble features of the planets, such as their relative sizes and orbits around the sun. R-simulations can be run using the scale model to predict or explain properties of the actual solar system. The model can represent the solar system because the model resembles (at a smaller scale) relevant states of the solar system. Since the evolution of the simulation process mimics the target process by its very construction, no separate representations of how the target process evolves given particular inputs are required—i.e., a theory of the target domain is not required for R-simulation. (See Goldman, 2006, however, for ways R-simulation and theorizing can work together).

Goldman (2006) defines the ST of mindreading in terms of R-simulation rather than CM-simulation. According to Goldman, we mindread by R-simulating a target person's mental states—i.e., we undergo mental processes which replicate or resemble the mental processes of the target— and use these interpersonal R-simulations to form beliefs about that person's mental states. The rival TT of mindreading, accordingly, involves CM-simulation. TT proposes that we mindread using a "theory" of human psychology, made up of generalizations about the relations between mental states, external stimuli, and behavior. This theorizing is a form of CM-simulation because it involves *descriptions* of how people's mental states and behavior change over time. In contrast, when engaging in interpersonal mental R-simulation, I entertain mental states that replicate/resemble the target's mental states, which I then attribute to the target. When engaging in mentalistic theorizing, the attributor represents the target's mental processes using mechanisms separate from those involved in actually having those mental processes herself—so the attributor's representations of the target's mental states do not replicate/resemble the target's mental states.

In summary, R-simulation requires the simulation process to replicate/resemble the simulated process, while CM-simulation need only represent the inputs, intervening states, and outputs of the target system. Applied to mindreading, ST says we mindread via interpersonal mental R-simulation, by activating our own psychological mechanisms in ways that replicate or resemble the mental processes of a target person. The theorizing about people's minds and behavior proposed by TT is, in contrast, is a form of CM-simulation, where the mental state representations use separate mechanisms that do not replicate those of the target person.

## A Sketch of Control Theory

I will now introduce some terminology from CT necessary for understanding Oztop et al.'s (2005) computational model. CT describes how a mechanism called a *controller* guides or controls the behavior of some system, called the *plant* or *controlled system* (Grush, 2004). In *feedback* control schemes, the controller receives input signals about the plant's goal state and its current state, and outputs control signals necessary for the plant to achieve the goal state. This mapping from goals to control signals is called the *inverse mapping*, and the controller often called an *inverse model*. When the plant implements a control signal, the controller receives sensory feedback about the new state of the plant. This is used to adjust the control signal so as to continue progress towards the goal state. This feedback loop continues until the system reaches its goal state.

More complex control schemes build on the idea of feedback control by introducing an additional way for the controller to obtain information about the plant: a duplicate control signal, or *efference copy*, is sent to a mechanism called a *forward model* whose function is to "model" the input-output function of the plant—i.e., the *forward mapping* from control signals to sensory signals about the plant's state. For a given control signal, the forward model produces an output signal about the plant's state that is identical/similar to the real feedback signal from the plant. The forward model thus can be described as predicting or representing the sensory feedback from the plant. Feedback from a forward model can reach the controller faster than real sensory feedback, allowing the controller to adjust its control signals sooner than would be possible in feedback control systems without forward models.

It is now commonly argued that our motor system uses such a control scheme (e.g., Wolpert, 1997). The body is controlled by the brain, and parts of the brain serve as inverse and forward models. Such a picture of motor control is the basis for Oztop et al.'s (2005) neurally plausible mechanism for visually guided motor control and action perception.

## A Control-Theoretic Model of Motor Control and Action Perception

Oztop et al.'s (2005) computational model has some notable limitations. It is restricted to the control of actions involving visible body parts (e.g., arm and hand movements). Accordingly, the architecture's extension to action perception can only understand these kinds of actions. But these (and other) simplifications, along with their successful computer simulations of the model's use, make it a useful example for evaluating the sense of "simulation" involved in CT-based architectures. I will first introduce the model's visual-feedback control architecture, and evaluate the sense of "simulation" involved here. Then I will do the same for the architecture's application to action perception.

### "Simulation" in Visually Guided Motor Control

To introduce how Oztop et al.'s model enables motor control via visual feedback, consider the following example. Suppose I want to grasp one of several objects located within reach in front of me. Given this goal, and visual information about the object's relative location and the

location of my limbs, my brain's controller produces a motor command that is relayed to my body. Accordingly, my arm begins to move toward the object, and my hand configuration adjusts so as to accommodate the object's shape. An efference copy of this motor command is sent to a forward model, which predicts the sensory signals—in this case, *visual* signals—that will be produced by this arm/hand movement. Sensory feedback from the body is slow compared to the forward model's predicted feedback. Accordingly, if the initial motor command is off track, the predicted feedback can permit faster adjustments than would be possible by waiting for real sensory feedback—i.e., waiting for the arm/hand to actually move, and visual feedback to travel from the retina to the controller. This complex feedback loop of motor commands followed by predicted and actual sensory feedback (which are used in different ways in different control architectures) is continued until I achieve my goal of grasping the object.

Oztop et al.'s (2005) visuomanual control mechanism (see Figure 1) follows this general description. Their model proposes that the parietal cortex, modulated by information about the actor's goal provided by the prefrontal cortex, extracts from visual input what is called the *control variable* ($X$): visual features about the body and environment relevant to the achievement of a given goal. For example, if the goal is to reach toward a particular location, $X$ is the distance between the index finger and the target location; if the goal is to grasp an object, $X$ will concern the distance between the relevant parts of the hand and the graspable parts of the target surface. The premotor cortex is the inverse model responsible for movement planning: it receives $X$ from the parietal cortex and the goal ($X_{des}$) from the prefrontal cortex, and outputs a desired change in body state ($\Delta\theta$), i.e., a control signal to be implemented by the body. The primary motor cortex and spinal cord use $\Delta\theta$ to generate the fine-grain motor commands necessary to move the body (forming the "dynamics control loop"). As the body moves, visual feedback is fed to the parietal cortex, giving the premotor cortex the new value of $X$.

This visual-feedback loop is supplemented by a "sensory forward model" which takes an efference copy of a motor command ($\Delta\theta$) and outputs $X_{pred}$, a prediction of the control variable to be computed by the parietal cortex at the next time step (after the body has moved). $X_{pred}$ is fed to the controller to inform subsequent motor commands. The forward model in this way can compensate for delays in visual feedback reaching the controller. Note that this forward model does not explicitly represent body states: it goes directly from desired changes in body state ($\Delta\theta$) to predicted sensory feedback ($X_{pred}$). In Grush's (2004) terminology, it is a "modal" forward model, rather than an "amodal" one containing explicit representations of the body from which mock sensory signals (of whatever sensory modality) can be generated. Oztop et al. discuss how such an amodal forward model could be used in their model, but do not include one.

Given this understanding of how the model works, we can ask: how is "simulation" involved? While Oztop et al. (2005) only mention simulation in the context of action perception, some of these same authors in other papers refer to motor control mechanisms as involving simulation. For example, Wolpert (1997) characterizes a forward model as "simulating the movement dynamics" of the body (p. 213). And Wolpert and Kawato (1998) generally describe an "internal model" (a forward or inverse model) as "a system which mimics the behavior of a natural process" (p. 1318). As Goldman (2006, pp. 216-217) argues, this language suggests that inverse and forward models perform R-simulation; but other common descriptions of CT's internal models—as "containing information" or "representing probabilities"—suggest the use of CM-simulation. Given this ambiguity, how are the internal models of Oztop et al.'s motor control mechanism best characterized?

Going beyond just the internal models, there are three main types of states to be found in the control mechanism (ignoring the dynamics control loop, and the visual cortex's sensory processing): (a) goals/intentions (e.g., in the prefrontal cortex); (b) control variables ($X$), i.e., visual representations of (desired, actual, or predicted) features of
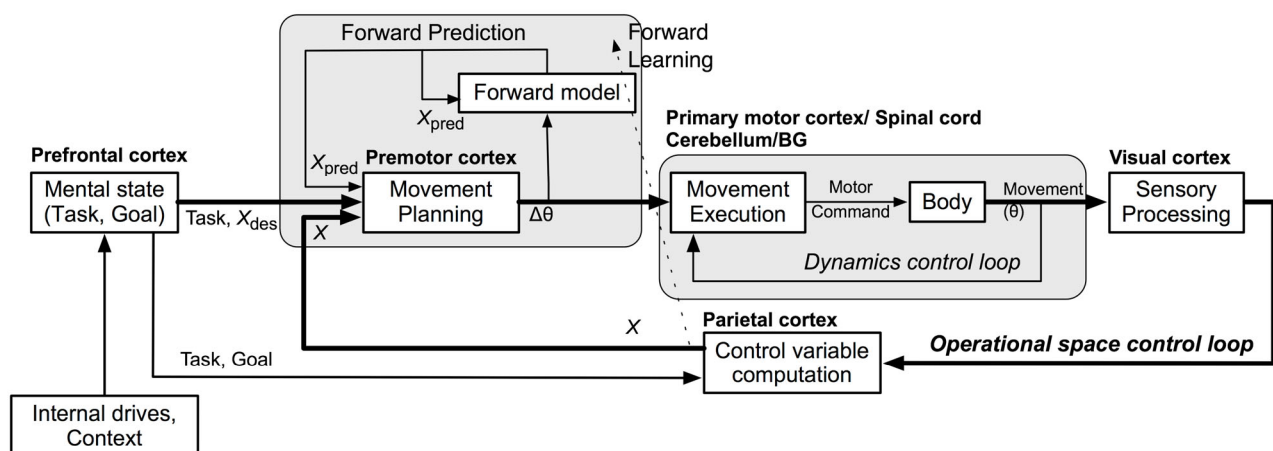


Figure 1: Oztop et al.'s (2005) visual feedback control mechanism (redrawn from Figure 1, p. 131).

the body/environment relevant to control; and (c) high-level motor commands (Δθ), i.e., desired changes in body states, which lead other parts of the brain to produce fine-grained motor commands. All three are arguably (indicative or imperative) representations of the body/environment. Goals and control variables are unproblematically representational. It is more controversial whether motor commands are representational, but let's assume they are (see Mandik, 2005). Since Goldman (2006, p. 132) allows that a simulation process can consist of a single state or event, we can ask whether these representational states are CM- or R-simulations. Remember that the distinctive feature of R-simulation is that the simulation *replicates or resembles* what it represents. Since the states of Oztop et al.'s visuomanual controller represent states of the *body/environment*, they clearly cannot count as R-simulations. As Goldman (2006, p. 217) argues, CM-simulation more aptly describes these representations.

Since the inverse and forward models are the parts explicitly identified as performing "simulation," I will examine them in more detail. Let's start with the inverse model. The premotor cortex produces motor commands so as to make the body attain its goal state. The authors do not say much about how the movement planner performs its input-output function. But regardless of how it works, note that "inverse model" is somewhat of a misnomer. The movement planner's function in motor control is not to "model" or represent anything outside the brain, but to produce control signals from information about the controlled system's current and desired state. That the input-output function of the movement planner is (ideally) the inverse of the input-output function of the body/environment does not seem enough to describe the controller as "simulating" the body/environment. But even if it did, there would be no role for R-simulation here.

What about the forward model, which takes in a control signal (Δθ), and outputs the predicted sensory effects of this movement ($X_{pred}$)? Assuming Δθ's are representations of desired changes in body state, the forward model's inputs and outputs *represent* (desired or predicted) features of the body/environment. The forward model must conform to the forward dynamics of the body/environment to produce accurate predictions. While this could be characterized as "mimicking" the forward dynamics of the body, this neural process clearly does not replicate the body/environment in the sense required of R-simulation. The forward model is thus better seen as CM-simulating the body/environment. Note that in characterizing CM-simulation, Goldman mainly had in mind theories, which explicitly represent states causally intervening between inputs and outputs. Since Oztop et al.'s sensory forward model does not explicitly represent the states intervening between input and output—i.e., states of the body—it might be characterized as a non-theoretical form of CM-simulation. The amodal forward models discussed above do, however, explicitly represent how motor commands affect body states, and the sensory signals produced by these body states. They are thus well

described as theories of the musculoskeletal system, representing generalizations of how motor commands affect body states, and body states affect sensory feedback.

There is, however, a sense in which the forward model uses R-simulation. The forward model's inputs and outputs do replicate/resemble other parts of the visuomanual control mechanism, other neural states of the same agent. The inputs are efferent *copies* of motor commands. And the forward model's outputs ($X_{pred}$) are supposed to replicate/resemble the real sensory feedback about the body computed by the parietal cortex ($X$). Accordingly, the forward model's inputs and outputs can be described as *intrapersonal* R-simulations, since the replicating and replicated states are "in the same individual mind" (Goldman, 2006, p. 37). Note, however, that its control-theoretic role as an "internal model" of the controlled system (the body) still does not involve R-simulation.

This analysis of Oztop et al.'s model leads to the following conclusions about CT-accounts of motor control. Goldman is right that the internal models responsible for motor control do not perform their representational functions by R-simulating the body's engagement with the environment. It is doubtful that inverse models are really "models" at all, but forward models are quite clearly representational, and can reasonably be described as CM-simulations of the body/environment. Finally, it can be argued that forward models R-simulate other brain activities (i.e., are intrapersonal R-simulations). But they do not R-simulate anything outside of the brain.

## "Simulation" in Action Perception

As depicted in Figure 2, Oztop et al. modify their visuomanual controller to enable action perception—i.e., understanding the goals or intentions behind the actions of (visually) perceived agents. The observer starts with some assumption about the actor's goal. While the nature of the "Estimated Mental State" box is not well specified, it is likely just a representation of the actor's goal. But the goal state found in the prefrontal cortex—the "Mental state (Task, Goal)" box—is more than a mere representation of a mental state: it is the goal state the observer would be in if she were to perform that action. This goal state is fed to the parietal cortex and the premotor cortex (movement planner).

The parietal cortex accordingly computes a control variable ($X_{observed}$) from visual information about the actor's observed action. $X_{observed}$ is thus the control variable the observer's parietal cortex would output if she were performing that action. The connection from parietal cortex to premotor cortex is inhibited. Instead, premotor cortex computes Δθ from information about the goal provided by prefrontal cortex ($X_{des}$) and the forward model's prediction of the control variable for that time step ($X_{pred}$). (This can only get off the ground by first initializing the forward model during a period of observing no movement.) The connection between premotor cortex and the areas responsible for motor execution is also blocked, so the observer does not actually move when premotor cortex

computes $\Delta\theta$. Instead, $\Delta\theta$ is fed only to the forward model, which predicts the control variable to be observed after implementing $\Delta\theta$ (i.e., $X_{pred}$). This production of predicted sensory signals from goal state estimates by the movement planner and forward model constitutes the "movement simulation." This movement simulation loop runs multiple times to create a sequence of predicted sensory signals.

A "difference" mechanism then compares $X_{pred}$ to $X_{observed}$, to determine whether the hypothesized goal produces sensory signals from the "simulated" movement ($X_{pred}$) that match the sensory feedback from observing the actor's actual movements ($X_{observed}$). If they match, the observer attributes the hypothesized goal to the actor. If there is a mismatch, an error signal is produced, leading to a change in the estimated goal state (the process labeled "mental state search"), driving another movement simulation. This "mental state inference loop" continues until a match is found, and that goal attributed to the actor.

In summary, Oztop et al.'s model enables action perception by producing an estimate of the actor's goal state, and "simulating" (a) the motor commands that would be produced to achieve this goal and (b) the sensory feedback from observing this movement. This predicted sensory feedback is tested against the real sensory feedback obtained from visually observing the actor. The model looks to be a use of interpersonal R-simulation Goldman (2006) calls a "generate-and-test" strategy: an observer generates hypotheses about the mental states responsible for some observed behavior, "then 'tests' (one or more) of these hypotheses by pretending to be in these states, feeding them into an appropriate psychological mechanism, and seeing whether the output matches the observed evidence. When a match is found…he attributes the hypothesized state or combination of states to the target" (p. 45). Note that the generate-and-test strategy is not a "pure" form of ST: the processes which generate the interpersonal mental R-simulations, and those which test these R-simulations against the observed evidence, are not themselves R-simulations. My tasks here will be to determine whether the

"movement simulation" identified by Oztop et al. indeed constitutes interpersonal R-simulation, and how to characterize the other parts of the mechanism.

Let's start with the parts added to the visuomanual controller especially for testing hypotheses about the actor's goal: the mental state estimate and the "difference" module. That neither mechanism has an analog in the actor helps us see that neither does any representing by R-simulation. The mental state estimate is a representation of the actor's goal separate from the observer's own goal states, and so does not play its representational role by replication/resemblance. It is thus a CM-simulation. The same holds for the "difference" module's comparison of predicted and observed sensory signals.

Continuing with the "hypothesis testing" process, consider the production of $X_{observed}$. The parietal cortex clearly performs the same function as during motor control: it computes control variables from visual input. It thus counts as a form of *intrapersonal* R-simulation. Note that this is a different sense of intrapersonal R-simulation than the one found earlier with efferent copies. There, one brain area replicated/resembled the activity of a different brain area. Here a single brain area designed for one function (motor control, its "online" function) is co-opted for use in a different cognitive activity (action perception, its "offline" function). Many characterizations of ST have considered such cases of intrapersonal R-simulation—where the "offline" operation of a mechanism is for mindreading— as criterial for simulating as opposed to theorizing about another person's mental states (e.g., using my decision-making mechanism "offline" to represent another person's decision-making). But on Goldman's account, *intrapersonal* R-simulation is fairly common, and *interpersonal* mental R-simulation additionally requires that the attributer's mechanism replicate/resemble what it is representing. Thus, being an intrapersonal R-simulation is not enough for the parietal cortex's activity to count as an interpersonal mental R-simulation. $X_{observed}$ represents the visual features of the actor's body/environment relevant to achieving the goal
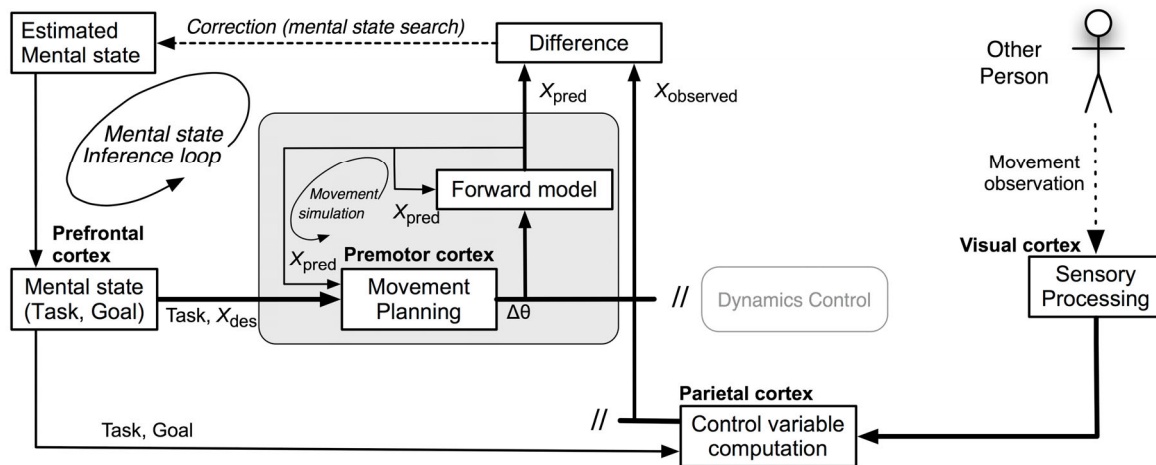


Figure 2: Oztop et al.'s (2005) "mental state inference" system for action perception (based on Figure 2, p. 133).

hypothesized to be the actor's. If the actor were really pursuing this goal, and using visual feedback to do so, her parietal cortex would be computing the same control variable. But the observer's parietal cortex does not perform its representational function by replicating the actor's parietal cortex activity. $X_{observed}$ represents the actor's body/environment, which it cannot replicate or resemble. It thus is not an interpersonal R-simulation.

I can now evaluate whether the "movement simulation" involves R-simulation. This sequence from the prefrontal cortex's goal state to the forward model's output of $X_{pred}$, obviously replicates the mental processes involved if the observer were herself acting—i.e., these processes are intrapersonal R-simulations. The question is whether these representations are interpersonal R-simulations, which requires that they represent analogous states in the actor's brain by attempting to replicate them. This is clearly the case for the prefrontal cortex and movement planner. These mechanisms represent the goal state representation, desired control variable, and motor command of the actor by replicating these states in the observer, rather than constructing descriptive representations of these states in separate mechanisms. But what about the forward model, which takes a replicated motor command and outputs a predicted control signal? If the actor's motor control system is like the observer's, she also has a forward model representing the forward dynamics of her body. But the observer does not represent predicted visual features of the actor by replicating the activity of the actor's forward model. Assuming her behavioral repertoire would not change, the actor could even stop using her forward model during motor control without affecting the role of the observer's forward model output in action perception—namely, to drive further R-simulations of the inverse model, and to be compared with $X_{observed}$ by the "difference" module. That the observer's forward model is a motor control mechanism co-opted for action perception and that the actor uses a similar mechanism do not necessitate that it is an interpersonal R-simulation. Instead, it is a CM-simulation of the human musculoskeletal system, in this case, the actor's. That the forward model first develops to represent the observer's own body does not detract from its being a CM-simulation when used to represent the actor's body. Oztop et al.'s "movement simulation loop" is thus a combination of R-simulation (by the prefrontal cortex and inverse model) and CM-simulation (by the forward model).

In summary, many of the mechanisms in Oztop et al.'s model of action perception are intrapersonal R-simulations: they replicate/resemble neural processes that occur when the observer is herself acting. But the notion of intrapersonal R-simulation does not distinguish between cases of interpersonal CM- and R-simulation. Just because a psychological process is activated in two different contexts (e.g., in motor control and in action perception), does not mean that its interpersonal use involves replicating the psychological processes of another person. With Oztop et al.'s model, not all the intrapersonal R-simulations are

interpersonal R-simulations. The processing stream from goal states in the prefrontal cortex to the output of motor commands by the movement planner indeed replicates these processes inside the observed actor. These processes constitute the "mental simulation" part of the generate-and-test strategy. But the parietal cortex's computation of control variables from visual observation of the actor, and the forward model's predictions of these sensory signals, are not interpersonal R-simulations. These representations are responsible for testing the accuracy of the movement R-simulations, rather than replicating anything inside the actor, and thus are CM-simulations. Thus, Oztop et al.'s description of their "movement simulation loop" is somewhat misleading, since it obscures the fact that both R-simulation and CM-simulation are involved.

## Conclusion

Interdisciplinary research in cognitive science often makes use of a host of conceptual frameworks. It is essential to determine whether a term common to multiple frameworks (here, "simulation") is being used in the same or different senses. Such conceptual issues are essential to accurately characterizing the phenomena at issue. For example, researchers studying the mechanisms of mindreading should attend to the distinction between intra- and interpersonal R-simulation. Discovering that a brain mechanism functions during both self- and other-oriented activities (e.g., acting and perceiving others' actions) is not enough to show that it accomplishes the latter via interpersonal R-simulation. The connection between "mirror neurons" and the ST of mindreading might be less direct than is generally assumed, if mirror neurons constitute forward models (Oztop et al., 2005) which perform CM- rather than R-simulation.

## References

Goldman, A. I. (2006). *Simulating minds: The philosophy, psychology, and neuroscience of mindreading.* Oxford, England: Oxford University Press.

Grush, R. (2004). The emulation theory of representation: Motor control, imagery, and perception. *Behavioral and Brain Sciences, 27*, 337–442.

Mandik, P. (2005). Action-oriented representation. In A. Brook & K. Akins (Eds.), *Cognition and the brain: The philosophy and neuroscience movement.* Cambridge, England: Cambridge University Press.

Oztop, E., Wolpert, D., & Kawato, M. (2005). Mental state inference using visual control parameters. *Brain Research Cognitive Brain Research*, *22*(2), 129–151.

Wolpert, D. M. (1997). Computational approaches to motor control. *Trends in the Cognitive Sciences*, *1*(6), 209–216.

Wolpert, D. M., Doya, K., & Kawato, M. (2003). A unifying computational framework for motor control and social interaction. *Philosophical Transactions of the Royal Society of London, Series B, 358*, 593–602.

Wolpert, D. M., & Kawato, M. (1998). Multiple paired forward and inverse models for motor control. *Neural Networks*, *11*(7–8), 1317–1329.