Hierarchies in Action and Motor Control

Sebo Uithol, Iris van Rooij, Harold Bekkering, and Pim Haselager

Abstract

■ In analyses of the motor system, two hierarchies are often posited: The first—the action hierarchy—is a decomposition of an action into subactions and sub-subactions. The second—the control hierarchy—is a postulated hierarchy in the neural control processes that are supposed to bring about the action. A general assumption in cognitive neuroscience is that these two hierarchies are internally consistent and provide complementary descriptions of neuronal control processes. In this article, we suggest that neither offers a complete explanation and that they cannot be reconciled in a logical or conceptually

coherent way. Furthermore, neither pays proper attention to the dynamics and temporal aspects of neural control processes. We will explore an alternative hierarchical organization in which causality is inherent in the dynamics over time. Specifically, high levels of the hierarchy encode more stable (goal-related) representations, whereas lower levels represent more transient (actions and motor acts) kinematics. If employed properly, a hierarchy based on this latter principle of temporal extension is not subject to the problems that plague the traditional accounts. ■

INTRODUCTION

In motor control, it is common to think of actions as hierarchically structured: A goal is served by an action, which, in turn, is served by multiple subactions. For example, when I want a glass of milk from the fridge, I have to get up from my chair, walk to the kitchen, open the door of the fridge, grasp the box of milk, and so on. I get up by means of placing my hands on the armrests, bending forward, stretching my legs, and pushing off. I place my hands on the armrest by means of stretching my arms, grasping the rests, and so forth (similar to Newell & Simon's, 1972, means-end structure in problem solving; see also Byrne & Russon, 1998). When the goal of getting a glass of milk is placed on top, and the other aspects of the action are arranged below it, a hierarchy appears. When going down the hierarchy, the tree gets wider (more elements on one level), although the elements become less abstract, down to the level of individual muscle movements.

A general assumption of cognitive science is that such action hierarchies are mirrored in the neural representation underlying them (Botvinick, 2008; Bechtel & Richardson, 1993). In other words, there are two hierarchies: an *action hierarchy* describing the action and a *control hierarchy* describing the neural processes that are presumed to bring the action about.¹ Cognitive scientists assume, either implicitly (Hamilton & Grafton, 2007) or explicitly (Botvinick, 2008), that these two hierarchies match. However, as Badre notes: "The fact that a task can be represented hierarchically does not require that the action system itself consist of structurally distinct processes" (Badre, 2008, p. 193);

so, this assumption should be subject to testing. But, whether these two hierarchies are identical is only partly an empirical matter. Before experiments to test this assumption can be designed, some important conceptual issues need to be addressed.

There are multiple ways to construct a hierarchy, but two hierarchical structures seem prevalent in the literature on action and motor control: One is a hierarchy based on constitutional or part-whole relations between the elements; the other is structured around a causal influence between the elements. When describing the action hierarchy, typically, a part-whole structure is presumed, whereas the control hierarchy is usually explained using a causal framework. However, we will show that these two structuring principles are in fact mutually exclusive, which suggests that the action hierarchy need not be similar to the control hierarchy. We will discuss empirical evidence that these two hierarchies are indeed dissimilar.

The remainder of this article is organized as follows. We will start by briefly elucidating the relation between actions and goals. Next, we will discuss the two main structuring principles of hierarchies in the motor domain and argue that they are incompatible and dissimilar. As an alternative account, we will discuss models that use different time scales for different control processes. In these models, structures can be found that can be seen as hierarchically structured but in a different and much more implicit form. This interpretation of a hierarchy is not subject to the problems that plague the first two options and might therefore be an interesting alternative for structuring elements in motor control. Understanding the nature of this hierarchical structure can guide empirical research into action control.

ACTIONS AND GOALS

The topmost level of a hierarchy in the motor domain is often labeled the “goal level” (Hamilton & Grafton, 2006), “desire level” (Grafton & Hamilton, 2007; Hamilton & Grafton, 2007), or “intention level” (Pezzulo, Butz, & Castelfranchi, 2008), but other labels, such as “super-ordinate action,” can be found as well (Humphreys & Forde, 1998). Below that, there is usually at least one level for “actions” (Hamilton & Grafton, 2006) or “subgoals” (Hamilton, 2009), and the bottom level is often labeled “movements” or “kinematics.” The exact labels of these levels may, of course, vary, as long as confusion is prevented. For reasons of clarity and consistency, we will call the elements on the highest level “goals,” the action features on lower levels “actions,” and the elements on the lowest level “motor acts,”² as can be seen in Figure 1.

The idea of a hierarchical structure in actions has been applied both to action execution and action observation or action understanding. The rationale behind this dual application is that there is evidence that the same brain structures are used for action generation and action observation (see the extensive body of literature on mirror neurons: Rizzolatti & Sinigaglia, 2010; Rizzolatti & Craighero, 2004; motor resonance: Uithol, van Rooij, Bekkering, & Haselager, 2011a; Fadiga, Fogassi, Pavesi, & Rizzolatti, 1995). Our analysis, however, is based mainly on claims about hierarchies in the execution of an action but might have consequences for action observation as well.

To get a better understanding of what is actually claimed when it is proposed that action production is structured hierarchically, it is useful to formulate an answer to two questions: (1) What makes one level higher than another (what is the variable on the vertical axis), and (2) what is the relation between features on different levels (what do

Figure 1. A typical action hierarchy with one goal level and multiple levels for actions and motor acts. Note that this hierarchy is far from complete; for every action, only a few subactions are shown.

Figure 2. A hierarchy, again simplified, structured according to the part-whole principle. The structure is practically identical to the hierarchies found in the literature on action representation (see Figure 1).

the lines between the hierarchy elements in Figure 1 portray)? By answering these two questions, we will be able to compare the two different accounts of hierarchical structures in the motor domain.

PART-WHOLE RELATIONS

We can interpret a hierarchy as portraying part-whole relations between the elements. Each level of the hierarchy comprises a set of subsystems, which are themselves composed of smaller units. For example, the action of “getting milk” consists of “walking to the fridge,” “opening the door,” “grasping the box of milk,” and so on. “Opening the fridge,” in turn, consists of “grasping the handle,” “pulling,” and so on (see Figure 2 for an example of a part-whole hierarchy). In such a hierarchy, “getting milk” does not exist apart from “walking to the fridge,” “opening the fridge,” and so forth; it is composed of these action features. In other words, when there is the right kind of reaching and opening and closing of the hand (and a milk box, of course), there is “grasping of the box of milk.” Likewise, when all the actions in the hierarchy are present, the goal of “getting milk” is present. In this case, the vertical axis denotes constitutive complexity. The higher up the axis, the more subparts in total a certain action element has. The lines in Figure 2 portray a “part of” relation.

Some important points need to be made in respect to a hierarchy based on a part-whole relationship between the elements. First, this hierarchy can be postulated independently of an underlying cognitive mechanism. It is a description of an action and a way of carving an action into smaller subactions and sub-subactions. The lower the hierarchical level, the more detailed the description is; the

higher the level, the more encompassing the element is. “Grasp handle” is just a label for the combination of “reach toward handle” and “full hand grip.” As such, it provides a description of the explanandum, not an explanation. Similarly, one can describe a human being as consisting a trunk, a head, two legs, and two arms. The head consists of eyes, ears, a nose, a mouth, and so forth. This description does not directly offer a mechanical explanation of the functioning of the human body; it describes the elements that need to be explained. This nature becomes evident when one tries to imagine how the postulated hierarchy could be refuted. It is hard to imagine empirical evidence that could show that “reaching” appears not to be a part of “grasping the milk box.” It seems that the kind of evidence that could refute this hierarchy would rather be conceptual in nature.

Next, the part-whole hierarchy does not allow causal influence between the elements, as that would mean that an element would be the cause of its own parts, and, in general, nothing can be the cause of its own parts (Craver & Bechtel, 2007; Lewis, 2000).³ Likewise, the head is not the cause of the eyes or nose. In terms of actions, this means that the reaching action cannot be the cause of the full hand grip but, also, that the goal of getting milk cannot be the cause of walking to the kitchen, which is at odd with most studies into goal-directed action. This suggests that the part-whole principle might not be the only principle at work in the general perception of a hierarchy in the motor domain.

Lastly, we have previously shown (Uithol et al., 2011a) that goals can be formulated as an action of a more abstract form (grasping a cup serves the goal drinking), a desired world state (grasping the cup to have a clean table), or an object (the cup is the goal of my grasping action). It is possible to construct a part-whole hierarchy only when goals, actions, and motor acts are of a similar nature, in this case, a type of action. Only goals formulated as a type of action have subparts that can be accommodated in a hierarchy. When goals are rendered as desired world states or objects, no relevant subparts of an action goal can be formulated and placed in a hierarchy. Objects of course have subparts (a cup has a handle, a saucer, etc.), but object parts have no place in an action hierarchy, as actions cannot be subparts of an object. The same goes for a desired world state: It has many (dissimilar) elements, such as objects and relations or properties, but they cannot be arranged in an action hierarchy. A part-whole hierarchy could be construed for a desired world state, but it would describe the world state, not the action needed to bring it about.

In summary, a hierarchy strictly based on a part-whole principle describes the action and its structure. No causal influence can be assumed between the different elements in the hierarchy. Consequently, a hierarchy strictly based on a part-whole principle may provide a characterization of an action but does not provide an explanation of actions or motor control. Also, a hierarchy of this type

allows only one interpretation of a goal, namely, a goal formulated as an action of a higher abstraction.

CAUSAL RELATIONS

An alternative principle to structure a hierarchy in the motor domain, not based on part-whole relations, is a causal hierarchy in which parts higher on the hierarchy are the cause of, or causally influence, parts lower on the hierarchy.⁴ The goal of getting a glass of milk activates a “get up” action, which activates a “stretch legs” and “bend trunk” action. In a causal hierarchy, higher-level elements can modulate the activity of lower-level mechanisms.

This structure differs from the part-whole structure in four important ways. First, the action features are not subparts of features higher up the hierarchy but necessarily exist independently of action elements higher in the hierarchy. It is important to realize that this renders the part-whole hierarchy and the causal hierarchy incompatible. In the part-whole hierarchy, the higher elements consist of the lower elements and, therefore, by definition, do not exist independently. In the causal hierarchy, the causal influence between the elements necessitates independent existence of the various elements.

Second, when goals exist independently of actions, it is no longer necessary that elements higher in the causal hierarchy are more complex than elements lower in the hierarchy. A simple element can just as well be the cause of a complex element. Indeed, goals and intentions are often posited to be discrete, constitutionally simple and propositional states (Pacherie, 2008; Haggard, 2005; see also Uithol, Burnston, & Haselager, submitted).

Third, possible interpretations of the notion of goals are no longer restricted to abstract action type of goals. The fact that parts need to be of a similar (ontological) nature as the whole entails that a part-whole hierarchy only allowed goals defined in terms of an action. This restriction drops out in a causal hierarchy so that goals formulated as a desired world state or an object can also be the cause of an action. Additionally, elements such as “affordances” (Gibson, 1977)—being a relation between an organism and an object—can now be accommodated.

Fourth, unlike the part-whole relation, the causal structuring principle does make claims about the underlying cognitive mechanisms. Effects and causes are assigned to different elements, and for these elements to have a physical reality, they must be assumed to be related to physical causes and effects, for instance, such as those that may hold in the brain.

To illustrate the nature of this hierarchy, let us assume that the goal of “getting milk” is the cause of “walking to the kitchen,” “opening the fridge,” and “grasping the milk.” When we want to add further detail to this hierarchy, for example, by further specifying “open fridge” into “full hand grip” and “pull,” we have to choose between simply replacing the element “open fridge” with this sequence of elements (Figure 3A) and adding an extra layer

Figure 3. Two different hierarchies structured according to the causal principle. In hierarchy (A) there is no extra control layer between “getting a glass of milk” and “full-hand grasp,” whereas in (B) “open fridge” exists as an independent causal unit.

below “open fridge” (Figure 3B). The difference is not a mere difference in visualization but actually corresponds to two different claims about the control of the action. In the latter situation, we postulate an extra control layer, which is ontologically independent of “full hand grasp” and “pull handle.” In this case, it is claimed that “opening the fridge” exists as a separate entity (a representation or a command), independent of the lower-level features.

In the causal hierarchy, the vertical axis denotes causal influence. Higher levels have causal influence on lower levels, but lower levels have no influence on higher levels. However, motor control is generally not believed to be instantiated by unidirectional downward causation. More realistic models of motor control implement feedback by means of reciprocal connections (Kilner, Friston, & Frith, 2007) and feedforward and error predictions (Friston, 2005; Haruno, Wolpert, & Kawato, 2001).

However, feedback between action elements on different levels is problematic to accommodate in a hierarchy structured around causal influence, as feedback is also a form of causal influence. If motor acts can also influence actions, and actions also influence goals, we seem to have lost the principled reason for placing goals on top and means at the bottom of the hierarchy. In other words, there seem to be no principles for placing one level below or above another level, which means a departure from one of the main characteristics of the control hierarchy: its top-down organization.

To make things causally even more complex and interconnected, in addition to the aforementioned interlevel causal influence, there is evidence for intralevel causal influence as well: Elements on a given level seem to influence each other. As an example, Cohen and Rosenbaum (2004) show a “hysteresis effect.” This effect shows that, during a grasping task, a previous grip location influences the location where an object is grasped next, even when this means that the well-known “end state comfort” principle (Rosenbaum & Jorgensen, 1992)—a presumable top-down process—has to be violated. As another example, Selen, Franklin, and Wolpert (2009) found that the “stiffness” used in pushing an object was not only an effect

of the characteristics of the object that was being pushed but also of the previous object. In other words, it mattered what the participant did before for how the task was executed.

There is also evidence that what you will do next influences on how you perform the current action or motor act. In speech articulation, this effect is known as “coarticulation” (Rosenbaum, 2009). When, for example, pronouncing “tulip,” the lips already round before pronouncing the “t” to correctly pronounce the “u,” but consequently, the “t” is pronounced slightly different.

When there seems to be mutual influence between elements on different levels as well as between elements on a single level, and we hold on to causality as the only principle for structuring the hierarchy, the image that emerges is more like a mesh with dynamically interconnected action features than a neat tree structure with an inherent top-down ordering of levels. In a tree with bidirectional causal influence, no unambiguous ordering of levels is implied by the causal relation alone.

To be clear, the conclusion of our analysis is not that the idea of a control hierarchy is in itself wrong. We have argued that, if such a hierarchy exists, then, it cannot be based on causal relations alone. Likewise, we do not wish to deny the existence of causal relations between the action elements, but framing the hierarchy entirely in terms of causal influence just does not seem to capture the complexity of influences present in the neural control of an action.

Still, we, as well as many other species, are capable of organizing our behavior in such ways that a predetermined goal is achieved. When I want a glass of milk, I usually have this goal before initiating action. Also, I usually succeed, regardless of a few obstacles on my path, and when necessary, I can adapt my behavior to unforeseen environmental demands and still succeed. This must mean that the goal of getting a glass of milk in Figure 3 has a dominance of some sort over the other action features. A clue on how this dominance could be achieved can be found in recent modeling work. We will discuss this in the Temporal Extension section. First, we will formulate the consequences of the

incompatibility explained above for cognitive research into motor control.

DIFFERENT HIERARCHIES FOR DIFFERENT PARTS OF THE EXPLANATION

Both the part-whole structure and the causal structure can be found in the literature on action representation and motor control. For example, Grafton and Hamilton (2007) provide much evidence for a form of distributed representation of an action in which different action elements are represented in different brain regions. They claim that this distributed nature of action representation provides evidence for a hierarchy in motor control. They note that “control hierarchies should be reflected by differences in those areas that are recruited for preparation and execution” (p. 599), suggesting a causal influence between the various elements. Later, however, they postulate an action hierarchy based on levels of complexity (p. 605), suggesting a part-whole structure.

In general, each of the hierarchies seems to have found its own niche within explanations of an action. When describing the action hierarchy, a hierarchy is often constructed on basis of the part-whole structure. The action is carved into subactions and sub-subactions, as explained above. On the other hand, when the control hierarchy is described, a causal structure is presumed. An overview of our conclusions thus far is presented in Table 1.

We have argued that the two structuring principles are not compatible. So, when the action hierarchy is supposed to be mirrored in the control hierarchy, a structuring principle that is applicable to both the hierarchies is needed. Unfortunately, neither the part-whole structure nor the causal structure seems to thrive outside its niche.

The causal structure makes little sense in the action hierarchy. We might be able to explain that my walking to the fridge is caused by the goal of getting milk, but it does not make sense to state that my leg swinging is caused by my walking, as that would entail that my walking could exist independently of leg swinging.

Applying a part-whole structure to a control hierarchy is equally problematic. First, as explained above, a part-whole hierarchy would not relate to a causal mechanism

but to a (complex) representation of an action at best. Second, when one is looking for a part-whole hierarchy in neural structures, one assumes that the structure in the content of the representation is mirrored in the structure of the vehicle of the representation, which means that one is looking for an action representation with a constituent structure (Fodor, 1975) or a microfeature structure (van Gelder, 1999). In this form of representation, the vehicle (i.e., the neural state that carries the information) has identifiable subparts, and content can be attributed to these subparts. Moreover, the content of the overall representation is dependent on the content of the subparts. So, in case of action representation, the goal representation should consist of subrepresentations that can be identified as actions. These subrepresentation again have subparts with identifiable content. For example, the representation of grasping the handle should consist of two identifiable representations: reaching toward the handle and a full hand grip. This strong restriction renders much of the available neural data insufficient to support a part-whole hierarchy, as not only do we have to find different representations for different subparts of an action, but these representation together also need to be correlated with the presence of a goal. So, for example, goal-sensitive mirror neurons in the macaque’s premotor cortex (Gallese, Fadiga, Fogassi, & Rizzolatti, 1996; Rizzolatti, Fadiga, Gallese, & Fogassi, 1996; Rizzolatti et al., 1987) cannot be accommodated in a control hierarchy based on a part-whole relation. The vehicle of this goal representation is simple in the sense that no functional subparts are known to date⁵ (Uithol, van Rooij, Bekkering, & Haselager, 2011b).

In all, the two structures are not compatible, and neither structure is transferable to the other side of the explanation. A direct consequence is that the control hierarchy and the action hierarchy need not match. Both the structure and set of elements of the two hierarchies can differ. Apparently, our intuition to divide an action into even smaller parts—our “folk motor control,” so to speak—might not be the best strategy for finding the neural correlates of action control.⁶ Indeed, Dennett warns us against the uncritical acceptance of a seemingly (intuitively) reasonable task description: “Marr’s more telling strategic point is that, if you have a seriously mistaken view about what the computational level description of your system is ..., your attempts to theorize at lower levels will be confounded by spurious artifactual puzzles. What Marr underestimates, however, is the extent to which computational level (or intentional stance) descriptions can also mislead the theorist who forgets just how idealized they are” (Dennett, 1989, p. 108). Instead, a constant interplay between gathering neural data and adapting the action hierarchy might be a more fruitful strategy.

Thus far, we have based our conclusion that the action hierarchy need not match the control hierarchy solely on conceptual grounds. In the next section, we will discuss

Table 1. The Two Types of Hierarchies and Their Properties

	<i>Action Hierarchy</i>	<i>Control Hierarchy</i>
Structuring principle	Part-Whole	Causality
Location	In the nervous system	In the action
Nature	Decomposition of explanandum	Mechanism

empirical evidence that there are in fact dissimilarities between these two hierarchies.

NEURAL EVIDENCE FOR TWO DIFFERENT HIERARCHIES

There are two ways in which the action hierarchy and the control hierarchy can be dissimilar: The control hierarchy can contain elements that are absent in the action hierarchy, and, vice versa, the action hierarchy can contain elements that are absent in the control hierarchy. There seems to be empirical evidence for both types of mismatches. To give an example of the first, Graziano and Aflalo (2007) stimulated the premotor areas of macaque monkeys for a relatively long duration (500–1000 msec). They were thereby able to evoke complex movement sequences to a certain end location, for instance, a sequence consisting of grasping, bringing to the mouth, turning the head toward the hand, and opening the mouth. Importantly, these movements were complex but “dumb”: When something blocked the trajectory of the bringing-to-the-mouth movement, the arm got stuck and did not move (Graziano, 2010, p. 461). These data seem to suggest that the behavioral repertoire of the monkey is represented by means of basic chunks and modifications to these chunks, such as target localization and adaptation to the trajectory when an object is blocking the pathway. However, a straightforward decomposition of the action into an action hierarchy would not automatically lead to these basic action chunks and, therefore, would not posit the additional modifying elements. This demonstrates that the control hierarchy contains elements that are absent in a straightforward action hierarchy.

Similarly, the most straightforward or intuitive decomposition of a grasping action is into the movements of individual fingers and the thumb. However, there is evidence that, at the neural side, the control of the grip is not decomposed into the movements of individual fingers but to a base posture with addition of refinements in finger and thumb position (Mason, Gomez, & Ebner, 2001). So, a straightforward decomposition of a precision grip grasping action would lead to an index finger and thumb movements as basic chunks, whereas the neural control hierarchy has a full hand grasp and suppression of three fingers as basic chunks. Again, our “folk” decomposition of an action seems not to correspond to the control hierarchy: The neural representation can contain elements that, at first sight, do not seem to be part of the action.

There seems to be neurological evidence for the opposite possibility as well: The control hierarchy can lack elements that do seem to be part of the action hierarchy. The literature on embodied,⁷ embedded cognition provides many examples of elements that can be considered part of an action but lack a neural correlate (see, for instance, Chiel & Beer, 1997). Clear examples can be found

in the human gait. Our gait is a complex orchestra of movements in many joints. The muscle activation responsible for a successful gait is hypothesized to be controlled by central pattern generators (Duysens & Van de Crommert, 1998). However, these neural patterns are not sufficient to generate a fluent and efficient gate. Passive components, such as muscle and tendon elasticity and inertia of the upper and lower leg, are of crucial importance (Whittington, Silder, Heiderscheit, & Thelen, 2008). In other words, some particular stages or parts of an action are not controlled by the neural patterns that activate muscles, but these stages are accomplished by “exploiting” regularities of the body, such as muscle and tendon elasticity, and the context, such as inertia and gravity, and are, in that sense, not centrally controlled but via self-organization. These important features of a normal gate are not part of the action representation but are, nevertheless, part of the action.

The problems outlined above suggest that, in their purest form, the two traditional principles for structuring a hierarchy might neither separately nor combined be the best candidates for a general theory on action representation. An interesting alternative for (or modification to) structuring the control hierarchy can be found in the temporal ordering of hierarchical elements or processes (Kiebel, Daunizeau, & Friston, 2008; Koechlin, Ody, & Kouneiher, 2003; Kelso, 1995). The fundamentals of such a hierarchy are best introduced by discussing a recent model in robotics (Yamashita & Tani, 2008). After this brief excursion, we will return to neuroscience and discuss Koechlin’s “cascade model” of neural control (Koechlin et al., 2003) that seems to be structured around the same principle.

TEMPORAL EXTENSION

Yamashita and Tani (2008) modeled a motor system of a robot without using what they call “local representations”: neural nodes dedicated to the representation of single action primitives in an explicitly represented hierarchical structure. Every 1 of the 180 units was connected to every other unit, including itself. The network was trained using backpropagation. They realized self-organization of a functional hierarchy through the use of two distinct types of neurons, each with different temporal properties. The first type of neuron was fast in the sense that its activity can change quickly. The second type of neuron was slow. They found that, after training, continuous sensorimotor flows are segmented into reusable motor primitives during repetitive execution of behavioral tasks. Moreover, these primitives could be flexibly integrated into new behavioral sequences. The model accomplishes this without setting up an explicit subgoal or function. In other words, without explicit instructions, representations of independent action elements emerge.

It is important for our analysis that the two types of neurons each developed a distinct activation profile. During

the execution of a repetitive motor task, repetitions of similar patterns were observed in activities of the fast context units. The activity in the slow units, in contrast, remained constant throughout the repetitious task. These results can be interpreted such that the fast units encoded reusable motor primitives that, because of their fast dynamics, were unable to preserve goal information over long trajectories. The slow context units, in contrast, encoded the switching between these primitives and, on account of their slow dynamics, could contribute to more stable goal representations. It is important to realize that the behavior of the robot was the result of the interplay of the different units and not of the slow units controlling the faster ones.

This interpretation could provide us another and less problematic structuring principle for a hierarchy: temporal extension. Elements higher on the hierarchy are represented longer or are more stable than lower ones. As such, they are able to influence an action for a longer time interval, thereby accounting for our capacity to structure behavior around a goal. In a way, this reverses the general reasoning: Elements are not more influential because they are higher in the hierarchy, but elements are higher in the hierarchy because they have more influence (on account of being more persistent).

Although it is related to causal influence, temporal extension is a different criterion for building a hierarchy. It is not assumed that the causal influence works in only one direction from goal to action—remember that every unit in the network Yamashita & Tani used was connected to every other unit. Nor is it assumed that the causal influence in one direction is bigger than in the reversed directed. The difference between the types of influence is a difference in temporal extension: Goals simply exert their influence longer than the actions or motor acts.

The control hierarchy structured on the basis of temporal extension is not committed to the direct causal influences as found in the causal hierarchy. This means that, although the overall structure—goals high in the hierarchy and action means low in the hierarchy—can be preserved, the hierarchy is much more implicit, as the orderly tree structure is lost. There is a simultaneous influence of a great many action features at an unbound number of levels.

The model built by Yamashita and Tani (2008) developed a functional hierarchy of only two layers, slow and fast, and the functional elements they found would still be located at the very bottom of the common action hierarchies. They suggest that “[t]he idea of functional hierarchy that self-organizes through multiple time scales may as such contribute to providing an explanation for puzzling observations of functional hierarchy in the absence of an anatomical hierarchical structure” (p. 13). Indeed, human action control seems to be hierarchically structured, as argued above, without a clear anatomical hierarchical structure (Miller & Cohen, 2001), so this relatively simple model illustrates a possible neural sub-

strate of an influential neurocognitive model on action control.

Koechlin, Basso, Pietrini, Panzer, and Grafman (1999) proposed a model in which different types of action control are located along a rostro-caudal axis in the lateral pFC. In their hierarchical model, four types of control are discerned (Koechlin & Summerfield, 2007; Koechlin et al., 2003). Sensory control, located at the caudal end of the axis, is involved in selecting motor actions. A bit more anterior, contextual control is involved in selecting pre-motor representations or stimulus–response associations. Next, episodic control is involved in selecting task sets or sets of consistent stimulus–response associations in the same context. Lastly, branching control, implemented in the rostral end of the axis (the anterior and frontopolar regions of pFC), involves controlling the activation of subepisodes nested in ongoing behavioral episodes.

The significance of the proposed model does not lie in the fact that exactly four different control layers are posited (it is, we believe, unlikely that human action control consists of fixed and integer number of control layers) but in the suggestion that different control processes operate on different time scales. When going from sensory control to branching control, the temporal extension of the types of control grows. Sensory control deals with selecting immediate movements—analogue to Yamashita and Tani’s fast neurons—and monitoring stimulus changes. The input to contextual control is already less transient. Episodic control deals with entire sets of association within one context, whereas branching control is involved in managing changes between different contexts. This means that these control processes can be accommodated in a hierarchy structured around stability or temporal extension.

Once this hierarchy is established, it is compelling to interpret Koechlin et al.’s finding in terms of a more traditional, causal control hierarchy, with an action goal originating in the higher control processes that is subsequently propagated to the lower types of control to evoke the appropriate action. By referring to their model as the “cascade model” and by emphasizing the downward modulation, Koechlin and colleagues are (perhaps unintentionally) feeding this compelling intuition, which is subsequently adopted by other researchers (Hamilton, 2009; Badre & D’Esposito, 2007).

However, the data do not support such an interpretation. Koechlin et al. (1999) show that, when more temporally extended forms of control are needed, anterior and frontopolar areas are activated in addition, not alternatively, suggesting that these control processes are not responsible for the control task by themselves but through interaction with the lower types of control, just like how all the units in Yamashita and Tani’s (2008) model contributed to the resulting behavior. This collective contribution is incompatible with the idea that goal-directed behavior is the result of higher layers propagating goal representations to lower layers. Goal-directed behavior emerges from the interaction between the different

types of processes, not from straightforward top-down modulation.

If we were to accept Koechlin's alternative hierarchy based on temporal extension but continued to interpret this control hierarchy as a top-down causal structure, we would not do justice to the complexity seemingly inherent in action control. The dynamic interaction between the various processes operating on different time scales is not captured by straightforward top-down causation. Additionally, interpreting the proposed hierarchy in terms of causal effects entails positing functionally discrete states that, through interaction, bring the action about. As extensively argued by Uithol et al. (submitted), such states seem inconsistent with existing data on action control, as well as conceptually incoherent, to propose that functionally discrete states with such causal effects lie at the basis of our actions. The informationally and dynamically complex control processes in the pFC are irreconcilable with the idea that discrete states are the primary cause of our action.

This insight could guide future research into action control. Instead of positing an anteriorly represented action goal and trying to locate the processes by which this representation is transformed to a motor program, the analysis above suggests that research into action control is better served by focusing on how goal-directed behavior emerges from the interaction between the different control layers. Which sensory input is used at which layer of control? How do lower control processes shape higher ones, and vice versa? Koechlin and colleagues made an important step in shifting this focus. This shift is hampered, however, if we allow the traditional views back in to shape our analysis.

An important theoretical advantage of an implicit hierarchy based on temporal extension is that it rids us from the rather artificial constraint that an action is associated with just one goal, present in both the causal and part-whole hierarchies. At every moment, one can be attributed many, maybe even an infinite number of, goals: to breathe, to read, to maintain homeostasis, to be a good scientist, to remain an upright posture, and so forth. Our behavior is the result of the interplay of this multitude of goals (McFarland, 1989; see also Uithol et al., submitted). These goals need not be represented on the higher layers in Koechlin's model but can also be an emergent result of the interaction of different control processes. To give a simple example, when swimming using a front crawl, a typical pattern of strokes and breathing is adopted. This pattern only makes sense when one realizes that two goals, to swim as fast as possible and to breathe, are pursued at the same time. Of course, we know about a swimmer's goal to breathe, and this goal is unlikely to be represented in one of Koechlin's control layers. In straightforward cognitive descriptions, we are inclined to leave it "out of the equation" and treat it as a boundary condition. But making a distinction between variables and boundary conditions in such an intuitive

and implicit manner might not be the best approach to a general theory on motor control. Although cognitive scientists generally have good reasons not to put an infinite number of goals in a model on action representation, to assume that the number of goals is always limited to only one might, in some cases, be overly restrictive.

This influence of multiple simultaneous goals cannot be easily accommodated in an explicit control hierarchy. An element can be caused or modulated by multiple goals at the same time. It might not always be clear what goals influence a lower element or to what extent. The result would be that the orderly tree-shaped hierarchy gets replaced by a dense mesh of interconnected action elements, which would seriously undermine the value of a hierarchy in explaining the realization of actions.

On the other hand, the more implicit hierarchy structured around is not committed to the postulation of a single, explicit goal nor a direct and univocal relation between the higher and lower elements. Therefore, the influence of multiple simultaneous goals does not undermine the hierarchical structure.

CONCLUSION

In theories on motor control, two hierarchies, the action hierarchy and the control hierarchy, are thought to match. We have presented both conceptual and empirical evidence suggesting that this assumption is unlikely to be true. We have shown that, implicitly, two structuring principles are used to construct a hierarchy but that neither structure (nor the [impossible] combination of the two structures) can provide an adequate framework for explaining actions and motor control. The action hierarchy, constructed using a part-whole hierarchy, is a description of the action that is to be explained but can be misleading in searching for a neural implementation of the action. The control hierarchy, constructed using causal relations, does not capture the complexity inherent in motor control. Our conclusion is not that motor control is not structured hierarchically at all but that the traditional accounts of an action hierarchy do not capture the complex and dynamic nature of motor control. Alternatively, dynamic accounts of motor control can be interpreted as hierarchical as well. In these models, elements that are represented longer and more stable are higher in the hierarchy. Although these alternative models are hierarchical in a much more implicit way and cannot straightforwardly be interpreted along the same lines of the more traditional accounts, they do not suffer from the conceptual and empirical issues discussed. Much work, both conceptual and empirical, is still needed to develop an implicit hierarchical structure around temporal extension to an insightful and coherent alternative to the current theories on action representation. Only if we approach the alternative hierarchy as a genuinely alternative structure and avoid straightforward causal interpretations based

on the traditional accounts, however, we can expect to find its true value.

Acknowledgments

We thank Dan Burnston and Emily Cross for constructive comments on an earlier draft of this article. This work was supported by a Donders internal graduation grant to H. B. and P. M., an NWO-VICI grant to H. B., and the EU-Project Joint Action Science and Technology (IST-FP6-003747).

Reprint requests should be sent to Sebo Uithol, Donders Institute for Brain, Cognition and Behaviour, Radboud University Nijmegen, PO Box 9104, 6500 HE Nijmegen, The Netherlands, or via e-mail: S.Uithol@donders.ru.nl.

Notes

1. To prevent confusion: In the literature on motor control, the notion “action hierarchy” is used for a hierarchical structure both in the action and in the neural control of the action. Here, we reserve the term for a hierarchical structure in the action or the behavior. Posited structures in the neural control of an action will be called “control hierarchy.”
2. When the ideomotor terminology is adopted, an action is a movement that serves a goal (Arbib & Rizzolatti, 1997). The elements in a hierarchy serve a goal by definition (otherwise, they could not be accommodated in the hierarchy), so the elements on the lowest level cannot be “mere” movements (i.e., not serving a goal), as they are sometimes referred to. Hence, we choose the term “motor act.”
3. Circular causality is a much debated concept within the dynamical system theory (Bakker, 2005; Lewis, 2005; Juarrero, 1999; Kelso, 1995; Port & van Gelder, 1995; Haken, Kelso, & Bunz, 1985) and means that elements on a lower level collectively contribute to a higher-level variable, which in turn modulates the behavior of elements at the lower level. As it is still highly contentious whether the downward causation (required for genuine circular causality) actually amounts to a causative force over and beyond the collective interactions of lower-level elements (Kim, 1993, 2000), we do not wish to pursue this issue here. More importantly, for our purposes, even if downward causation in this strong sense would exist, the claim still is not that the collective variable would actually cause its own parts (i.e., their existence as parts) but, instead, that it would causally constrain their behavior and would therefore fall under the second principle to structure a hierarchy (see below).
4. Although this is generally true for action hierarchies, in the perceptual hierarchy, the order is reversed: Features low in the hierarchy, such as lines and colors, are thought to be the cause of higher-level features, such as objects (Felleman & Van Essen, 1991; Hubel & Wiesel, 1959).
5. Features such as spiking frequency or phase could play a functional role in the representational capacities of a neuron. To our knowledge, no study investigated these properties of mirror neurons.
6. The fact that the action hierarchy might not map (perfectly) onto the control hierarchy also has interesting consequences for theories on action understanding by means of motor resonance or mirror neurons. In these theories, it is assumed that, when observing action, the same neural structures are recruited as when executing an action (Uithol et al., 2011a). But when the features of action control do not match the action features we distinguish in an observed action, the nature of the “shared representations” (de Vignemont & Haggard, 2008) needs to be subjected to further research.

7. The notion of “embodiment” is used for various forms of dependency on a body. In cognitive science, it can refer to something as modest as the activation of the motor cortex (de Vignemont & Haggard, 2008), whereas usually in philosophy, a more radical mutual dependency between body and brain in the generation of behavior is meant (Haselager, van Dijk, & van Rooij, 2008; van Dijk, Kerkhofs, van Rooij, & Haselager, 2008; Clark, 1997; see Ziemke & Kirsh, 2003, for an overview of the various interpretations). Here, we use the more radical interpretation.

REFERENCES

- Arbib, M. A., & Rizzolatti, G. (1997). Neural expectations: A possible evolutionary path from manual skills to language. *Communication and Cognition*, 29, 393–424.
- Badre, D. (2008). Cognitive control, hierarchy, and the rostro-caudal organization of the frontal lobes. *Trends in Cognitive Sciences*, 12, 193–200.
- Badre, D., & D’Esposito, M. (2007). Functional magnetic resonance imaging evidence for a hierarchical organization of the prefrontal cortex. *Journal of Cognitive Neuroscience*, 19, 2082–2099.
- Bakker, B. (2005). The concept of circular causality should be discarded. *Behavioral and Brain Sciences*, 28, 195–196.
- Bechtel, W., & Richardson, R. C. (1993). *Discovering complexity*. Princeton, NJ: Princeton University Press.
- Botvinick, M. (2008). Hierarchical models of behavior and prefrontal function. *Trends in Cognitive Sciences*, 12, 201–208.
- Byrne, R., & Russon, A. (1998). Learning by imitation: A hierarchical approach. *Behavioral and Brain Sciences*, 21, 667–684.
- Chiel, H., & Beer, R. (1997). The brain has a body: Adaptive behavior emerges from interactions of nervous system, body and environment. *Trends in Neurosciences*, 20, 553–556.
- Clark, A. (1997). *Being there: Putting brain, body, and world together again*. Cambridge, MA: MIT Press.
- Cohen, R. G., & Rosenbaum, D. A. (2004). Where grasps are made reveals how grasps are planned: Generation and recall of motor plans. *Experimental Brain Research*, 157, 486–495.
- Craver, C., & Bechtel, W. (2007). Top-down causation without top-down causes. *Biology and Philosophy*, 22, 547–563.
- de Vignemont, F., & Haggard, P. (2008). Action observation and execution: What is shared? *Social Neuroscience*, 3, 421–433.
- Dennett, D. C. (1989). Cognitive ethology: Hunting for bargains or a wild goose chase? In A. Montefiore & C. Noble (Eds.), *Goals, no-goals, and own goals* (pp. 101–116). London: Unwin Hyman.
- Duysens, J., & Van de Crommert, H. (1998). Neural control of locomotion: Part 1. The central pattern generator from cats to humans. *Gait & Posture*, 7, 131–141.
- Fadiga, L., Fogassi, L., Pavesi, G., & Rizzolatti, G. (1995). Motor facilitation during action observation: A magnetic stimulation study. *Journal of Neurophysiology*, 73, 2608–2611.
- Felleman, D., & Van Essen, D. (1991). Distributed hierarchical processing in the primate cerebral cortex. *Cerebral Cortex*, 1, 1–47.
- Fodor, J. (1975). *The language of thought*. New York: Crowell.
- Friston, K. (2005). A theory of cortical responses. *Philosophical Transactions of the Royal Society of London, Series B, Biological Sciences*, 360, 815.
- Gallese, V., Fadiga, L., Fogassi, L., & Rizzolatti, G. (1996). Action recognition in the premotor cortex. *Brain*, 119, 593–610.
- Gibson, J. J. (1977). A theory of affordances. In R. Shaw & J. Bransford (Eds.), *Perceiving, acting, and knowing: Toward an ecological psychology* (pp. 67–82). Hillsdale, NJ: Lawrence Erlbaum.

- Grafton, S., & Hamilton, A. (2007). Evidence for a distributed hierarchy of action representation in the brain. *Human Movement Science*, 26, 590–616.
- Graziano, M. (2010). Ethologically relevant movements mapped onto the motor cortex. In A. Ghazanfar & M. Platt (Eds.), *Primate neuroethology* (pp. 454–470). New York: Oxford University Press.
- Graziano, M., & Aflalo, T. (2007). Mapping behavioral repertoire onto the cortex. *Neuron*, 56, 239–251.
- Haggard, P. (2005). Conscious intention and motor cognition. *Trends in Cognitive Sciences*, 9, 290–295.
- Haken, H., Kelso, J. A. S., & Bunz, H. (1985). A theoretical model of phase transitions in human hand movements. *Biological Cybernetics*, 51, 347–356.
- Hamilton, A. (2009). Research review: Goals, intentions and mental states: Challenges for theories of autism. *Journal of Child Psychology and Psychiatry*, 50, 881–892.
- Hamilton, A., & Grafton, S. (2006). Goal representation in human anterior intraparietal sulcus. *Journal of Neuroscience*, 26, 1133–1137.
- Hamilton, A., & Grafton, S. (2007). The motor hierarchy: From kinematics to goals and intentions. In P. Haggard, Y. Rossetti, & M. Kawato (Eds.), *Attention & performance 22. Sensorimotor foundations of higher cognition attention and performance* (pp. 381–408). Oxford: Oxford University Press.
- Haruno, M., Wolpert, D., & Kawato, M. (2001). MOSAIC model for sensorimotor learning and control. *Neural Computation*, 13, 2201–2220.
- Haselager, P., van Dijk, J., & van Rooij, I. (2008). A lazy brain? Embodied embedded cognition and neuroscience. In P. Calvo & T. Gomila (Eds.), *Handbook of cognitive science. An embodied approach* (pp. 273–290). Oxford: Elsevier.
- Hubel, D., & Wiesel, T. (1959). Receptive fields of single neurones in the cat's striate cortex. *Journal of Physiology*, 148, 574–591.
- Humphreys, G., & Forde, E. (1998). Disordered action schema and action disorganisation syndrome. *Cognitive Neuropsychology*, 15, 771–812.
- Juarrero, A. (1999). *Dynamics in action: Intentional behavior as a complex system*. Cambridge, MA: MIT Press.
- Kelso, J. A. S. (1995). *Dynamic patterns: The self-organization of brain and behavior*. Cambridge, MA: MIT Press.
- Kiebel, S. J., Daunizeau, J., & Friston, K. J. (2008). A hierarchy of time-scales and the brain. *Plos Computational Biology*, 4, e1000209.
- Kilner, J., Friston, K., & Frith, C. (2007). Predictive coding: An account of the mirror neuron system. *Cognitive Processing*, 8, 159–166.
- Kim, J. (1993). *Supervenience and mind: Selected philosophical essays*. Cambridge: Cambridge University Press.
- Kim, J. (2000). *Mind in a physical world*. Cambridge, MA: MIT Press.
- Koechlin, E., Basso, G., Pietrini, P., Panzer, S., & Grafman, J. (1999). The role of the anterior prefrontal cortex in human cognition. *Nature*, 399, 148–151.
- Koechlin, E., Ody, C., & Kouneiher, F. (2003). The architecture of cognitive control in the human prefrontal cortex. *Science*, 302, 1181–1185.
- Koechlin, E., & Summerfield, C. (2007). An information theoretical approach to prefrontal executive function. *Trends in Cognitive Sciences*, 11, 229–235.
- Lewis, D. (2000). Causation as influence. *The Journal of Philosophy*, 97, 182–197.
- Lewis, M. D. (2005). Bridging emotion theory and neurobiology through dynamic systems modeling. *Behavioral and Brain Sciences*, 28, 169–245.
- Mason, C., Gomez, J., & Ebner, T. (2001). Hand synergies during reach-to-grasp. *Journal of Neurophysiology*, 86, 2896.
- McFarland, D. (1989). Goals, no-goals and own goals: A debate on goal-directed and intentional behaviour. In A. Montefiore & D. Noble (Eds.), *Goals, no-goals and own goals* (pp. 39–57). London: Unwin Hyman.
- Miller, E. K., & Cohen, J. (2001). An integrative theory of prefrontal cortex function. *Annual Review of Neuroscience*, 24, 167–202.
- Newell, A., & Simon, H. A. (1972). *Human problem solving*. Englewood Cliffs, NJ: Prentice Hall.
- Pacherie, E. (2008). The phenomenology of action: A conceptual framework. *Cognition*, 107, 179–217.
- Pezzulo, G., Butz, M. V., & Castelfranchi, C. (2008). The anticipatory approach: Definitions and taxonomies. In G. Pezzulo, M. V. Butz, C. Castelfranchi, & R. Falcone (Eds.), *The challenge of anticipation. A unifying framework for the analysis and design of artificial cognitive systems* (pp. 23–43). Berlin: Springer-Verlag.
- Port, R. F., & van Gelder, T. (1995). *Mind as motion: Explorations in the dynamics of cognition*. Cambridge, MA: MIT Press.
- 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, 131–142.
- Rizzolatti, G., Gentilucci, M., Fogassi, L., Luppino, G., Matelli, M., & Ponzoni-Maggi, S. (1987). Neurons related to goal-directed motor acts in inferior area 6 of the macaque monkey. *Experimental Brain Research*, 67, 220–224.
- Rizzolatti, G., & Sinigaglia, C. (2010). The functional role of the parieto-frontal mirror circuit: Interpretations and misinterpretations. *Nature Reviews Neuroscience*, 11, 264–274.
- Rosenbaum, D. A. (2009). *Human motor control* (2nd ed.). San Diego, CA: Academic Press.
- Rosenbaum, D. A., & Jorgensen, M. J. (1992). Planning macroscopic aspects of manual control. *Human Movement Science*, 11, 61–69.
- Selen, L. P. J., Franklin, D. W., & Wolpert, D. M. (2009). Impedance control reduces instability that arises from motor noise. *Journal of Neuroscience*, 29, 12606–12616.
- Uithol, S., Burnston, D., & Haselager, P. (submitted). Is there a neural correlate of intentions?
- Uithol, S., van Rooij, I., Bekkering, H., & Haselager, P. (2011a). Understanding motor resonance. *Social Neuroscience*, 6, 388–397.
- Uithol, S., van Rooij, I., Bekkering, H., & Haselager, P. (2011b). What do mirror neurons mirror? *Philosophical Psychology*, 24, 607–623.
- van Dijk, J., Kerkhofs, R., van Rooij, I., & Haselager, P. (2008). Can there be such a thing as embodied embedded cognitive neuroscience? *Theory & Psychology*, 18, 297.
- van Gelder, T. (1999). Distributed versus local representation. In R. Wilson & F. Keil (Eds.), *The MIT Encyclopedia of Cognitive Sciences* (pp. 236–238). Cambridge, MA: MIT Press.
- Whittington, B., Silder, A., Heiderscheit, B., & Thelen, D. (2008). The contribution of passive-elastic mechanisms to lower extremity joint kinetics during human walking. *Gait & Posture*, 27, 628–634.
- Yamashita, Y., & Tani, J. (2008). Emergence of functional hierarchy in a multiple timescale neural network model: A humanoid robot experiment. *Plos Computational Biology*, 4, 1–18.
- Ziemke, T., & Kirsh, A. (2003). What's that thing called embodiment? In R. Alterman & D. Kirsh (Eds.), *Proceedings of the 25th Annual Conference of the Cognitive Science Society* (pp. 1134–1139). Mahwah, NJ: Lawrence Erlbaum.

This article has been cited by:

1. K. Richard Ridderinkhof. 2014. Neurocognitive mechanisms of perception–action coordination: A review and theoretical integration. *Neuroscience & Biobehavioral Reviews* **46**, 3–29. [[CrossRef](#)]
2. Fausto Caruana, Sebo Uithol, Gaetano Cantalupo, Ivana Sartori, Giorgio Lo Russo, Pietro Avanzini. 2014. How action selection can be embodied: intracranial gamma band recording shows response competition during the Eriksen flankers test. *Frontiers in Human Neuroscience* **8**. . [[CrossRef](#)]
3. Michiel van Elk, Hein van Schie, Harold Bekkering. 2014. Action semantics: A unifying conceptual framework for the selective use of multimodal and modality-specific object knowledge. *Physics of Life Reviews* **11**, 220–250. [[CrossRef](#)]
4. Bernhard Hommel. 2014. The role of action semantics in action control. *Physics of Life Reviews* **11**, 259–260. [[CrossRef](#)]
5. Sebo Uithol, Daniel C. Burnston, Pim Haselager. 2014. Why we may not find intentions in the brain. *Neuropsychologia* **56**, 129–139. [[CrossRef](#)]
6. Roy de Kleijn, George Kachergis, Bernhard Hommel. 2014. Everyday robotic action: lessons from human action control. *Frontiers in Neurorobotics* **8**. . [[CrossRef](#)]
7. Sasha Ondobaka, Harold Bekkering. 2013. Conceptual and perceptuo-motor action control and action recognition. *Cortex* **49**, 2966–2967. [[CrossRef](#)]
8. Sebo Uithol, Markus Paulus. 2013. What do infants understand of others' action? A theoretical account of early social cognition. *Psychological Research*. [[CrossRef](#)]
9. Nikolaas N. Oosterhof, Steven P. Tipper, Paul E. Downing. 2013. Crossmodal and action-specific: neuroimaging the human mirror neuron system. *Trends in Cognitive Sciences* **17**, 311–318. [[CrossRef](#)]
10. Mark Blokpoel, Johan Kwisthout, Theo P. van der Weide, Todd Wareham, Iris van Rooij. 2013. A computational-level explanation of the speed of goal inference. *Journal of Mathematical Psychology* **57**, 117–133. [[CrossRef](#)]
11. Matthew Michael Botvinick. 2012. Hierarchical reinforcement learning and decision making. *Current Opinion in Neurobiology* **22**, 956–962. [[CrossRef](#)]