APHASIOLOGY, 2006, 20 (9/10/11), 1112–1124

Psychology Press
Taylor & Francis Group

Gestural knowledge evoked by objects as part of conceptual representations

Daniel N. Bub and Michael E. J. Masson

University of Victoria, BC, Canada

Background: Theories of embodied knowledge argue that the representation and recruitment of motor processes may be important for deriving the meaning of many linguistic and perceptual elements.
Aims: We examined the conditions under which gestural knowledge associated with manipulable objects is evoked.
Methods & Procedures: A priming paradigm was used in which an object was presented in advance of a photograph of a hand gesture that participants were to mimic. On related trials, the target gesture was the same as the gesture typically used to interact with the object prime. On unrelated trials, the target gesture was not related to the object. In another set of experiments, a Stroop-like paradigm was used in which participants learned to produce manual responses to colour cues. After training, coloured photographs of manipulable objects were presented. The colour-cued gesture was either one typically used with the object or was unrelated to it.
Outcomes & Results: In the priming experiments, response latencies were shorter in the related condition, but only when participants also made an identification response to the object prime. In the Stroop experiments, interference effects indicated that gestures to colour were affected by gestural knowledge associated with the object.
Conclusions: These results indicate that conceptual representations of manipulable objects include specific forms of gestural knowledge that are automatically evoked when observers attend to an object.

Sentences often describe actions dealing with manipulable objects. For example, consider the sentence *Before the interview, Mary hastily applied some lipstick*. One view of sentence comprehension might take the meaning of this sentence to simply require an abstract understanding of the generic function that lipstick entails. Thus, the essential meaning of the sentence might rest on the knowledge that Mary is changing the colour of her lips by means of a cosmetic. But, clearly, the implications of a sentence may depend on more precise knowledge of how the action described is carried out. Suppose we are informed that Mary manipulated the lipstick by holding the cylinder in a clenched fist. Given our conventional understanding of how lipstick is used, it is immediately apparent that there is something anomalous about Mary or

Address correspondence to: Daniel Bub, Department of Psychology, University of Victoria, P.O. Box 3050 STN CSC, Victoria BC V8W 3P5, Canada. E-mail: dbub@uvic.ca

The authors made equal contributions to the work reported here. This research was supported by discovery grants from the Natural Sciences and Engineering Research Council of Canada to Michael Masson and to Daniel Bub, and by a grant from the Perceptual Expertise Network, which is funded by the James S. McDonnell Foundation. We are grateful to Marnie Jedynak for conducting the experiments.

http://www.psypress.com/aphasiology DOI: 10.1080/02687030600741667

at least about her present situation. We can also be sure that the result of her cosmetic effort is likely to be less than ideal.

The meaning of sentences involving actions on objects may therefore include either explicit or tacit information about how objects are actually manipulated. Indeed, recent accounts of language comprehension suggest that understanding manipulable objects or words referring to such objects requires a mental simulation of action. The knowledge of how a tube of lipstick is manipulated and used constitutes an essential part of the meaning of the concept "lipstick". Theories of embodied knowledge argue for an even more central role of action; the representation and recruitment of motor processes may be important for deriving the meaning of many linguistic elements, and the embodiment of meaning through action extends into abstract domains, including the comprehension and use of metaphor (Barsalou, Simmons, Barbey, & Wilson, 2003; Gallese & Lakoff, 2005; Glenberg & Kaschak, 2002).

There may be important functional relationships between motor representations and language, but how do we distinguish between different possibilities? At present we lack an adequate experimental approach that would reveal the dynamic evocation of actions when objects, words, and sentences are processed for meaning. Consider a number of important questions for which no answers are available at present: Are motor representations evoked by objects even when there is no accompanying intention to carry out an action on the object by the observer? If yes, then what is the nature of these representations? What role do they play in various language tasks? Do words evoke motor representations and under what circumstances? Do these motor representations differ from those evoked by objects? Objects inherently afford multiple actions; for example, we manually interact with a common object like a pocket calculator in several different ways when using it. Which of these actions, if any, are represented as part of the meaning of the object or word?

Functional imaging studies remain an ambiguous source of evidence, although widely cited as support for the claim that manipulable objects, or words and sentences referring to such objects, automatically recruit motor processing. Some experiments (e.g., Chao & Martin, 2000; Creem-Regehr & Lee, 2005) have indicated that passive viewing of tools is sufficient to evoke a range of specific cortical responses associated with motor processes. However, other findings suggest that visual objects do not invariably evoke motoric activation, but that such activation is task dependent. For example, Gerlach, Law, and Paulson (2002) showed premotor cortex involvement in a categorisation task (natural vs manmade), but not in object decisions (real vs non-real). Devlin et al. (2002) reported a meta-analysis of seven studies that used positron emission tomography to examine specific activation patterns for man-made objects, especially tools, in relation to other object classes (e.g., fruits and vegetables). They found evidence for activation in left posterior temporal regions that was specific to tools, but only when participants engaged in naming or semantic classification tasks, not during passive viewing. Clearly, the relationship between participants' task orientation to objects and the kind of premotor representations evoked remains an issue.

In this article we describe the progress we have made in developing experimental methods to reveal the evocation of specific hand actions in response to pictured objects or written words. For example, part of knowing what actions to carry out with a pocket calculator includes the depressing (poking) of keys. A crucial goal in evaluating the claim that lexical meaning includes access to action representations is to have ways of measuring the dynamic evocation of specific hand postures in a variety of tasks involving objects, words, and sentences.

Before describing our research, we summarise previous work in other laboratories that has gone some way towards establishing that certain aspects of hand movement are automatically recruited by objects or their names. We point out the limits of inferences allowed by this work. We then outline the logic of our own experimental approach, and we discuss what we have learned so far on the evocation of specific hand actions in relation to the meaning of manipulable objects and words that refer to them.

A number of previous studies have examined whether words or manipulable objects evoke the representation of hand actions. The basic design of these experiments requires participants to produce a manual response that is potentially influenced by knowledge of how one interacts with an object that is in view. Tucker and Ellis (1998), for example, had participants press a response button with the right or the left hand to classify the orientation of manipulable objects as upright or inverted. Each object was seen in profile and had a handle on its left or right side (e.g., a teapot). When the response hand was aligned with the handle of the object, participants were faster, implying the automatic evocation of a grasp response that affected the generation of the task-defined button press. Evidence that words can affect the generation of actions was provided by Glover, Rosenbaum, Graham, and Dixon (2004), who showed that at early stages of grasping a target block, the aperture between thumb and forefinger was affected by the presence of a word referring to a small or large manipulable object (e.g., apple or grape).

Although these studies demonstrate the evocation of some components of manual action (e.g., left versus right hand and finger aperture), neither addresses the central question of whether objects or words evoke specific hand postures (e.g., a poke gesture to a calculator) and under what circumstances. More recent studies by Tucker and Ellis (2001, 2004) have shown that semantic judgements about graspable objects were faster when the response action to signal the decision (a power or a precision grip) was compatible with the hand action usually associated with the target object. These results go some way towards establishing that specific hand postures are evoked by objects under certain task conditions. We note, however, that a wide range of hand actions in addition to prehensile grasps is typically needed to interact with objects, especially when using manmade objects for their intended purpose. Moreover, we wish to develop a more flexible task than one that requires manual responses driven by decisions that depend on explicit identification of an object. Thus, we present experiments that require manual responses to arbitrary cues and we examine how these responses are altered in the presence of manipulable objects. In addition, we measure a variety of possible hand gestures potentially evoked by objects including, but not limited to, prehensile grips.

EXPERIMENTS 1–3: PRIMING FUNCTIONAL GESTURES

Our approach to measuring the evocation of motor affordances is based on the development of a task that requires participants to initiate a pantomime gesture in response to a cue that defines some specific hand posture. These cues are relatively transparent; they are actual photographs of a hand in a posture that participants are asked to mimic (e.g., a hand with forefinger extended as a poke gesture). We measure the time to initiate each cued action by detecting when the response hand lifts away from a depressed response button. On each trial, we present an object either in advance of or simultaneously with a hand cue. The set of objects and associated hand actions are shown in Figure 1. Our method is analogous to a standard priming paradigm in the psycholinguistic literature, but instead of word recognition as the target task, we use the generation of a pantomime cued by the depiction of hand posture. The relationship between the object displayed as a prime and the hand cue is as follows. On *related* trials, the action denoted by the hand cue matched the action typically associated with the function of the object shown as the prime, as illustrated in Figure 1. On *unrelated* trials, the hand cue was paired with one of the other objects in the set, resulting in an unrelated object–hand pairing. If merely viewing an object without making a response to it elicits pertinent gestural knowledge, then participants should be able to generate actions to the cue more readily on related trials. Alternatively, it may not be the case that passive viewing is sufficient to elicit actions to objects. At least some evidence from the neuroimaging literature implies that semantic processing objects is required to generate motorically based cortical activity (Devlin et al., 2002). We therefore compared priming effects in two situations: participants passively viewed the object and responded to the hand cue or they were instructed to name the object after first responding to the hand cue.

In each of the experiments in this series, participants were shown greyscale photographs of objects and hands. The full set of items is shown in Figure 1. On each trial an object, serving as a prime, and a hand cue were presented. The participant's

Figure 1. The set of eight objects and corresponding hand gestures used in Experiments 1–3. Each gesture depicts the typical manual action applied when using an object for its intended purpose.

task was to mimic the posture formed by the hand cue. The critical trials began with the participant's dominant hand resting on a response key. To make a response, the participant lifted his or her hand from the key and used the dominant hand to mimic the posture shown in the cue. Latency to initiate this pantomime response was recorded automatically when the response was released. To ensure smooth responding, participants began a test session with a series of practice trials in which they repeatedly mimicked each of the hand gestures represented by the set of hand cues but in the absence of any prime stimulus.

In Experiment 1, participants passively viewed the object prime and no task requirement was associated with that stimulus. The prime was in view for 105 ms or 750 ms, then was replaced by the hand cue. The object prime on each trial was either related or unrelated to the hand posture that was to be made on that trial. For example, the calculator was a related prime for the poke gesture, whereas any of the other objects (e.g., hand saw) were unrelated primes. We tested 24 healthy undergraduate students. The manipulation of prime duration (105 vs 750 ms) did not interact with the prime relatedness factor in any of the experiments in this series, so we present the data collapsed across prime duration. Figure 2 shows the mean response latency in Experiment 1 as a function of prime relatedness. It is clear from the figure that no priming effect was obtained in this experiment. Mere passive viewing of the prime objects apparently failed to elicit gestural knowledge of the type that could influence production of a manual response.

In Experiment 2 we sought to have participants engage with the object primes by requiring them to report the name of the prime after they had mimicked the cued gesture on each trial. By having participants identify the prime object, we expected that they would need to recruit conceptual knowledge associated with the object. If this knowledge includes representations for gestural interactions with the object, then we should find shorter response latencies when the object and target gesture are related than when they are unrelated. The procedure in Experiment 2 was the same as in Experiment 1, except that after pantomiming the target gesture, participants also named the prime object. The mean response latency, shown in Figure 2, revealed

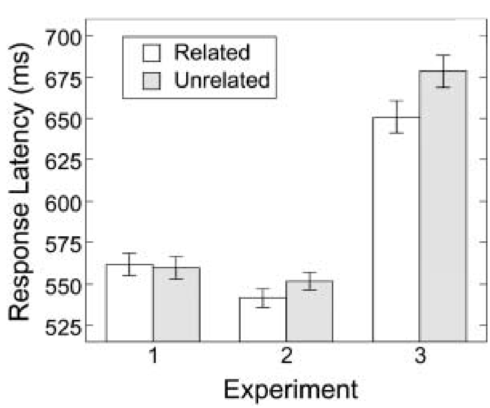

Figure 2. Mean response latency in Experiments 1–3 as a function of prime type. A related object prime represented an object to which the target gesture would typically be applied; unrelated object–prime/gesture–target pairs were created by reassigning objects and gestures. Error bars represent 95% within-participant confidence intervals (Loftus & Masson, 1994; Masson & Loftus, 2003).

a clear priming effect of about 11 ms ($n = 20$, $p < .01$). This result supports our expectation that identifying an object would potentially recruit gestural knowledge.

We attempted to produce a more robust priming effect in Experiment 3. Here, we modified the displays so that the prime object and the hand cue were presented simultaneously, side by side. As in Experiment 2, the task was to first pantomime the cued gesture, then to report the name of the object. By presenting the object and hand cue simultaneously, while requiring the participant to make a selective response to each stimulus, we expected to enhance the interaction between object identification processes and gestural knowledge. The mean response latency for pantomiming the indicated gestures is shown in Figure 2. Overall latencies in this experiment are longer because participants had to encode the object identity while viewing the hand cue and before initiating the manual response—the display was erased as soon as the participant's hand was lifted from the response key. The priming effect was magnified to 28 ms ($n = 21$, $p < .001$) in this experiment, relative to Experiment 2.

The evidence indicates, then, that merely looking at an object like a calculator is not enough to evoke a precise enough representation of hand action to prime the execution of a pantomime corresponding to the gesture normally made when using the object. Rather, objects generate representations of actions when classified and labelled. Our results cast some doubt on over-interpretations of functional imaging experiments that show evocation of motor-related activity when participants passively view objects. Either such activity occurs because objects are displayed long enough to encourage elaborative processing (e.g., speculation on how an object is used) or the motor representations observed in these experiments do not include the detailed parameters of action that led to priming in our task.

When participants are asked to name manipulated objects, however, specific hand actions are potentiated. Representations of hand actions are therefore evoked by the meaning of objects. A sentence like "Calculators are used for adding numbers" should not be taken as a literal index of the contents of semantic memory. Rather, such statements also reflect the fact that the knowledge of adding with a calculator includes elements of motor actions carried out on the physical device (Allport, 1985). Classifying and naming an object like a calculator involves access to this knowledge, yielding evocation of the corresponding representation of hand actions.

EXPERIMENTS 4–7: DISTINGUISHING GESTURES OF FORM AND FUNCTION

Objects normally elicit many kinds of manual actions. When using a calculator, for example, we often begin by picking it up and moving it to a convenient location for operating. A thumbtack requires that one pick it up with a pinch gesture before pressing it into a surface with an extended thumb. The prior history of an object is associated with multiple gestures. We can assume that frequently used actions are all represented as part of object knowledge. This assumption implies that hand actions for manipulable objects include actions corresponding to both the function of an object and its form. In some cases, form- and function-based actions are identical or nearly so. For example, we pick up and use a glass in virtually the same way. Other objects, though, have distinct actions for using and for picking up or moving them. A stapler is typically used to attach pages of paper by depressing its top component using a flat palm, but it is picked up using an inverted open grasp. We refer to form-based actions as *volumetric*, to reflect their sensitivity to details of weight distribution and shape. Actions based on conventional use of the object we refer to as *functional* gestures. This distinction is similar to one made by Johnson and Grafton (2003), who classify object-based actions as "acting on" and "acting with". The former, they assume, occurs when objects are grasped without a specific purpose, whereas the latter relies on knowledge of tool use. However, we do not consider volumetric gestures to be devoid of intentions or goals. Rather, they are executed for purposes such as picking up or moving, and in the case of familiar objects are influenced by memory for prior experience in much the same way as functional gestures.

In considering the possibility that the representations of both functional and volumetric gestures are part of the meaning of an object, we note that both types of gesture are ineluctably part of our prior motor experience with the object. If the knowledge of an object includes this prior history, then there is no reason to assume that functional gestural knowledge alone is maintained at the expense of volumetric knowledge. Indeed, functional gestures depend in crucial ways on knowledge of an object's shape and size encapsulated in volumetric constraints. The keys on a calculator are spatially arrayed within the confines of a particular volume, which determines their separation and therefore the positioning of the finger movements during calculation. We present a series of experiments, therefore, establishing that action representations evoked by objects include volumetric as well as functional gestures, and that both are part of the meaning of an object. After describing this evidence, we discuss the implications of our findings for the interpretation of motor-based activity elicited by objects in functional imaging studies and we consider the possible role that representations of volumetric and functional actions may play in the processing of words and sentences.

This series of experiments used a paradigm introduced by Bub, Masson, and Bukach (2003). The method was designed to provide evidence of the automatic nature of the recruitment of gestural knowledge when viewing an object, and is based on a variant of the Stroop colour-word interference task (Stroop, 1935). In the Stroop task, participants are instructed name the colour in which a stimulus is printed. If the stimulus is an incongruent colour word (e.g., *red* printed in green), participants take much longer to name the colour than if the stimulus is a neutral form such as a row of Xs or a congruent colour name (e.g., *red* printed in red). The long response latencies found in the incongruent condition are generally thought to indicate that participants automatically process word identity even though doing so can only interfere with their ability to efficiently perform the colour-naming task.

In our variant of the Stroop task, participants were first trained to associate each of four gestures with a unique colour. We used four functional gestures and four volumetric gestures in these experiments, testing different groups of participants with each set of four gestures. Figure 3 shows two examples of each type of gesture and an object for which each gesture would be appropriate. After learning the set of colour–gesture associations, participants were then given a series of trials on which they were to respond as quickly as possible to a colour stimulus by making the associated gesture. On these trials, colour was carried by an object (e.g., a green calculator). Colour–object pairs were arranged so that on congruent trials the target gesture was relevant to or congruent with the object (e.g., a green calculator shown to a participant who had learned to respond to green with a poke gesture). On incongruent trials, however, the object's colour was associated with one of the other gestures in the set, which always was an inappropriate gesture for that object (e.g., a

Figure 3. Example objects and gestures used in Experiments 4–7. The top row presents two functional gestures and an associated object for each. These gestures are the ones typically made when using the object for its intended purpose. The bottom row shows two volumetric gestures and an associated object for each. Volumetric gestures are the ones typically made when picking up an object, rather than using it.

red calculator shown to a participant who had learned to respond to red with a trigger gesture).

In this situation, participants are not required to make any response to the objects, but they must examine them to determine the colour in which they appear. However, participants should not attempt to recruit gestural knowledge about the object if they wish to avoid conflict on incongruent trials. But if participants are unable to control the evocation of gestural knowledge when viewing objects (because it occurs automatically), then we can expect that on incongruent trials participants will be slowed in making their responses, relative to congruent trials. For now, we leave aside the question of how the type of object viewing required in this task is different from the passive viewing assumed to operate in Experiment 1, but we will return to this issue in the General Discussion. In addition, one might ask whether participants were induced to recruit gestural knowledge about objects in these experiments because we included congruent trials, where gestural knowledge associated with the object would benefit production of the required response. This issue was addressed across Experiments 4–6 by varying the proportion of congruent trials included in the experiment.

In Experiment 4, we used congruent object–gesture pairs on 50% of the trials and incongruent pairs on the other 50%. In Experiment 5, we reduced the proportion of congruent trials to 25%. Finally, in Experiment 6, we dispensed with congruent trials and introduced a set of "neutral" objects that were large, non-manipulable entities (e.g., boat, airplane, sofa). In this latter case, there could be no motivation for participants to recruit gestural knowledge associated with objects because doing so could serve only to impede generation of the colour-cued gesture. Mean response latency for each of these experiments is shown in Figure 4. It can be seen that in nearly all cases, there was a clear disadvantage for responding in the incongruent

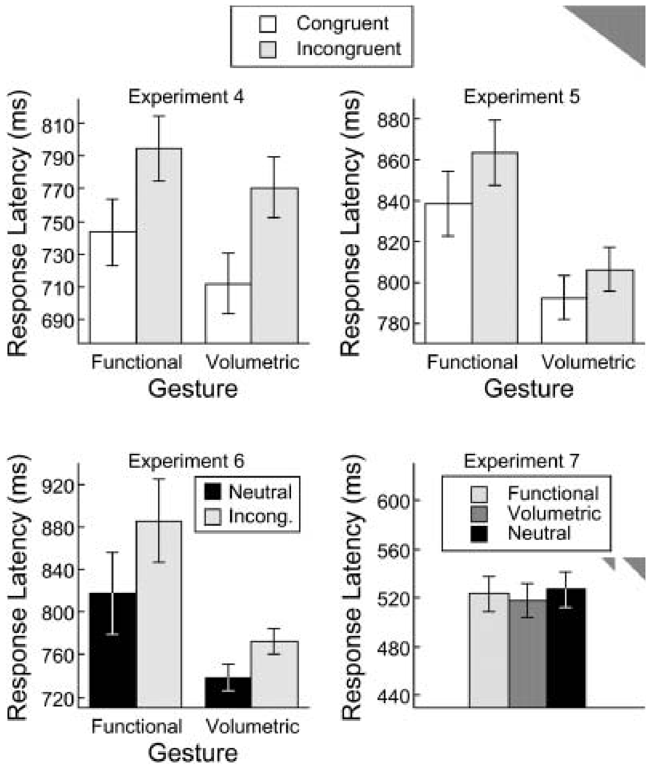

Figure 4. Mean response latency in Experiments 4–7. The data for Experiments 4–6 show gesture latency as a function of object–gesture relationship and type of gesture (functional or volumetric). The data for Experiment 7 show colour naming latency in the control experiment as a function of the type of object used to carry colour. Error bars represent 95% within-participant confidence intervals.

condition. The only exception is for volumetric gestures in Experiment 5, which produced a trend for longer durations in the incongruent condition. The incongruency effect in Experiment 6 is especially important because it shows that gestural knowledge seems to be evoked even when participants had no reason to strategically call upon their gestural knowledge of the objects they were shown, and indeed had reason to try to suppress such knowledge. This result speaks strongly in favour of the automatic nature of the evocation of gestural knowledge.[1]

One might argue that the reason for the incongruency effect seen in Experiment 6 is that participants could more easily discriminate colours when they were carried by the neutral objects than when the gesture-relevant objects were presented. After all, we could not counterbalance assignment of objects to congruency condition in Experiment 6, as we had done in the other two experiments. To dispense with this possibility, we had a group of participants perform a speeded colour-naming task in

[1] A reviewer pointed out that the trigger hand gesture physically resembles its associated object (the water pistol) and wondered whether this similarity might have played a major role in the congruency effects. This concern pertains only to the data from the functional gesture condition because the trigger gesture and the water pistol were only used there. We reanalysed the data for the functional condition in each experiment, excluding trials involving this object. A congruency effect emerged in each case even when this item was removed.

Experiment 7. Participants were shown the eight objects we used for functional gestures, the eight objects used for volumetric gestures, and the eight neutral objects. Each object appeared multiple times in each of the four colours used in Experiments 4–6 and the task was to name the colour as quickly as possible. Mean colour-naming latency is shown in Figure 4. It is clear that the three object sets did not differ with respect to colour-naming latency, indicating that colours were equally discriminable regardless of which set of objects carried them.

GENERAL DISCUSSION

The act of identifying manipulable objects evokes gestural representations associated with an object's function, and also with the object's form. These representations are part of object knowledge, but there are important constraints on the kind of the task context that elicits them. It is not the case that passive viewing alone is sufficient; pantomiming actions in response to a photograph of a hand denoting a particular posture was not affected by the accompanying image of an object unless participants were also instructed to name the object after carrying out the pantomime. Yet, surprisingly, an alternative experimental procedure yielded strong evidence that objects automatically evoke gestural representations. If participants pantomimed an action using colour as a cue, and the colour was a surface feature of an object, then both functional and volumetric gestures were elicited and interfered with production of the pantomime when there was a mismatch between it and the gestures conventionally associated with the object.

The paradox between these two outcomes—objects fail to elicit motor affordances in one context but readily do so in another—can easily be resolved if we consider that participants' use of colour as a cue required them to manually respond to an attribute of the depicted object (its colour) without also responding to its form. This is a very unusual filtering task: We are generally never in a situation demanding one kind of action to a surface property of an object, like colour, that is different from actions we normally execute when picking up or using it. Under these circumstances, the intention to act on the colour would not be easily segregated from the more habitual intention to carry out an action on the object, and motor representations afforded by the object are automatically evoked. When the cue to pantomime is a hand shape instead of an arbitrary colour, the response is to a depicted shape (a hand in a particular posture) that is completely distinct from the accompanying object. In this situation, no effect of the object is seen on pantomiming the hand cue, unless the participant also identifies and labels the object. We conclude that visual processing of manipulable objects does not invariably yield precise motor affordances. Gestural knowledge about form and function is evoked if participants attend to the meaning of an object (e.g., when asked to name it), but without intentions to act on the object, passive viewing alone does not automatically elicit gestural representations.

Nevertheless, the evidence we have obtained suggests that objects can readily generate motor representations in semantic tasks, and these representations concern both the function of an object and its volumetric properties. The fact that part of the manipulation knowledge associated with an object like a calculator includes parameters of hand actions executed when lifting the object seems surprising. We generally consider that familiarity with the meaning of an object entails knowing how to manually interact with it according to its conventional function, but this knowledge inevitably requires also knowing how to shape the hand when moving or lifting the object. Indeed, there is good reason on neuropsychological grounds to assume that familiar objects evoke stored knowledge of volumetric manual action as part of their meaning. Jeannerod, Decety, and Michel (1994) reported a case of optic ataxia in which hand configurations for grasping novel objects like a cylindrical tube were severely impaired, but manual interactions with familiar objects like a tube of lipstick were relatively preserved. Knowledge of a familiar object must include a representation of the volumetric gestures associated with shape-based properties, and this knowledge can modulate action even when cortical damage disrupts manual grasping of novel objects.

The fact that both functional and volumetric representations are triggered by objects is of considerable interest in regard to the interpretation of patterns of activation observed in neuroimaging experiments. Cortical areas known to mediate motor function are invoked when participants carry out tasks with tools and other manipulable objects, including tasks that do not require explicit consideration of manual actions (Devlin et al., 2002). It is generally assumed that this activation concerns manipulation knowledge dealing with the function of the object (e.g., Chao & Martin, 2000). Our evidence that hand actions pertaining to object shape, independent of function, are a crucial part of manipulation knowledge raises an important question. Specifically, are regions of activation associated with tools indeed indicative of functional knowledge, or do they represent both function and form? The finding that fruits and vegetables can yield activation of motor cortex (Gerlach et al., 2002) suggests that the representation of shape-based grasping is an important potential component of the observed patterns of activation.

How central are functional and volumetric gestural representations to the meaning of object concepts? Our experimental approach can be adapted to an examination of language tasks in order to determine the conditions generating manual representations evoked by words and sentences referring to manipulable objects. For example, written words instead of depicted objects can be used to carry colours that participants have learned to associate with particular gestures. Colour–word pairings can be arranged so that the gesture cued by the colour matches or mismatches the gestural knowledge associated with the manipulable object denoted by the word. For matching word–colour pairs, the relationship between the colour and the word can be volumetric or functional. Surprisingly, words show exactly the same evocation of functional and volumetric gestural knowledge as picture of objects (Bub, Masson, & Cree, 2006). This evidence strongly implies that gestural knowledge is part of the meaning of words and confirms a recent result reported by Myung, Blumstein, and Sedivy (in press). These authors found that making word–nonword decisions was facilitated on positive trials when the target word was preceded by another word that shared the same kind of manual action (e.g., typewriter, piano). In addition, Zwaan and Taylor (2006) showed that during comprehension of sentences referring to actions involving hand rotation, participants invoked motor representations that interacted with their perceptual judgements of a rotating image.

If gestural knowledge is recruited in processing the meaning of words, then the dynamic role of sentence context on motor representations raises some interesting questions. We have already clear evidence that the word concept "calculator" includes functional as well as volumetric gestural representations. Consider, however, a sentence like "John gave the calculator to Mary". The meaning of this sentence implies an action in which a calculator is being handed from one person to another. What kind of motor representations would be evoked?

Clues to the dynamics of motor representations evoked by sentences can be obtained by considering neurophysiological and computational accounts of contextual influences on action selection in primates (Fagg & Arbib, 1998; Sakata, Taira, Kusunoki, Murata, & Tanaka, 1997). In these accounts, the anterior-inferior parietal lobe (AIP) receives input from ventral pathways and generates multiple grasps afforded by an object. Grasp selection requires input to AIP from the inferior premotor cortex, based on contextual information that includes task demands and intentions. Irrelevant affordances are then suppressed, leaving only the affordance corresponding to the selected grasp.

In the human, we propose a similar dynamic resolution of affordances driven by factors such as sentence context. If understanding a sentence requires a mental simulation of action, then for a sentence like *John gave the calculator to Mary*, multiple affordances evoked by a calculator's shape and function would become available in AIP. The parameters for the relevant hand posture (open grasp) implied by the sentence would emerge over time and parameters for irrelevant gestures, such as poke, would dissipate. A sentence referring to function (*John added the numbers*) should generate the reverse pattern.

We return to the question of the causal role that gestural knowledge plays in language comprehension. The strongest claim is that understanding the meaning of words and sentences often or even invariably requires embodied action. The consistent effects of volumetric and functional gestural representations that we have described would then arise inevitably as a result of the operations required for language understanding. In a recent communication to us, John Marshall indicated that he was sceptical of this idea. We, who wish simply to follow the dictates of empirical evidence, suggest the following rejoinder. If a sentence that makes no appeal to visual or manual interactions with an object (e.g., *Jane remembered the calculator*) still evokes gestural knowledge, then we need to consider carefully the linguistic processes that generate motor representations and the possibility that these representations are an integral part of conceptual knowledge.

REFERENCES

Allport, D. A. (1985). Distributed memory, modular systems and dysphasia. In S. K. Newman & R. Epstein (Eds.), *Current perspectives in dysphasia*. Edinburgh, UK: Churchill Livingstone.

Barsalou, L. W., Simmons, W. K., Barbey, A. K., & Wilson, C. D. (2003). Grounding conceptual knowledge in modality-specific systems. *Trends in Cognitive Sciences*, *7*, 84–91.

Bub, D. N., Masson, M. E. J., & Bukach, C. M. (2003). Gesturing and naming: The use of functional knowledge in object identification. *Psychological Science*, *14*, 467–472.

Bub, D. N., Masson, M. E. J., & Cree, G. S. (2006). *Evocation of functional and volumetric gestural knowledge by objects and words*. Manuscript submitted for publication.

Chao, L. L., & Martin, A. (2000). Representation of manipulable man-made objects in the dorsal stream. *NeuroImage*, *12*, 478–484.

Creem-Regehr, S. H., & Lee, J. N. (2005). Neural representations of graspable objects: Are tools special? *Cognitive Brain Research*, *22*, 457–469.

Devlin, J. T., Moore, C. J., Mummery, C. J., Gorno-Tempini, M. L., Phillips, J. A., & Noppeney, U. et al. (2002). Anatomic constraints on cognitive theories of category specificity. *NeuroImage*, *15*, 675–685.

Fagg, A. H., & Arbib, M. A. (1998). Modeling parietal–premotor interactions in primate control of grasping. *Neural Networks*, *11*, 1277–1303.

Gallese, V., & Lakoff, G. (2005). The brain's concepts: The role of the sensory-motor system in conceptual knowledge. *Cognitive Neuropsychology*, *22*, 455–479.

Gerlach, C., Law, I., & Paulson, O. B. (2002). When action turns into words: Activation of motor-based knowledge during categorization of manipulable objects. *Journal of Cognitive Neuroscience*, *14*, 1230–1239.

Glenberg, A. M., & Kaschak, M. P. (2002). Grounding language in action. *Psychonomic Bulletin & Review*, *9*, 558–565.

Glover, S., Rosenbaum, D. A., Graham, J., & Dixon, P. (2004). Grasping the meaning of words. *Experimental Brain Research*, *154*, 103–108.

Jeannerod, M., Decety, J., & Michel, F. (1994). Impairment of grasping movements following a bilateral posterior parietal lesion. *Neuropsychologia*, *32*, 369–380.

Johnson, S. H., & Grafton, S. T. (2003). From "acting on" to "acting with": The functional anatomy of object-oriented action schemata. In C. Prablanc, D. Pelisson, & Y. Rossetti (Eds.), *Neural control of space coding and action production. Progress in brain research, Vol. 142* (pp. 127–139). Amsterdam: Elsevier.

Loftus, G. R., & Masson, M. E. J. (1994). Using confidence intervals in within-subject designs. *Psychonomic Bulletin & Review*, *1*, 476–490.

Masson, M. E. J., & Loftus, G. R. (2003). Using confidence intervals for graphically-based data interpretation. *Canadian Journal of Experimental Psychology*, *57*, 203–220.

Myung, J-Y., Blumstein, S. E., & Sedivy, J. C. (in press). Playing on the typewriter, typing on the piano: Manipulation knowledge of objects. *Cognition*.

Sakata, H., Taira, M., Kusunoki, M., Murata, A., & Tanaka, Y. (1997). The TINS Lecture – The parietal association cortex in depth perception and visual control of hand action. *Trends in Neurosciences*, *20*(8), 350–357.

Stroop, J. R. (1935). Studies of interference in serial verbal reactions. *Journal of Experimental Psychology*, *18*, 643–662.

Tucker, M., & Ellis, R. (1998). On the relations between seen objects and components of potential actions. *Journal of Experimental Psychology: Human Perception and Performance*, *24*, 830–846.

Tucker, M., & Ellis, R. (2001). The potentiation of grasp types during visual object categorization. *Visual Cognition*, *8*, 769–800.

Tucker, M., & Ellis, R. (2004). Action priming by briefly presented objects. *Acta Psychologica*, *116*, 185–203.

Zwaan, R. A., & Taylor, L. J. (2006). Seeing, acting, understanding: Motor resonance in language comprehension. *Journal of Experimental Psychology: General*, *135*, 1–11.