

on the trial immediately following a reward and then increasing steadily on subsequent nonrewarded trials.

The experiments consisted of recording from single DA neurons while the monkey performed the tasks. The analyses focused on how the DA responses changed as a function of the number of trials since the previous reward (“PRN”). To understand why, recall that DA responses are thought to encode a reward prediction error that is, in turn, used to make better predictions. A trial without a reward causes a certain reward prediction error. Thus, the more trials in a row without a reward, the more the prediction changes. By studying how the prediction error changes under these conditions, they could infer how the prediction is computed.

Results from the first task were consistent with the TD model. The DA responses—representing prediction error—increased steadily with PRN. As described by TD learning, for each unrewarded trial, the prediction was updated so that next time, no reward was slightly more strongly predicted than last time (Ivan Rodriguez did not get a big contract when he was a free agent last year, so he should have expected less money this year). Thus, when a reward did come, it was slightly more surprising and caused a larger DA response than last time. Conversely, when no reward came, it was slightly less surprising and caused a less inhibited DA response than last time.

Results from the second task were not consistent with the TD model, instead showing that context matters. For this task, an increasing PRN corresponded to an increasing probability that a reward would be delivered. Accordingly, DA responses decreased steadily with increasing PRN. These decreasing responses imply a decreasing prediction error that must therefore be based on a prediction of increasingly probable reward. Note that if context was ignored, as in the TD model, the responses would have increased with PRN, by the same logic as for the first task.

The results imply that our understanding of the computations performed by DA neurons is at best incomplete. This conclusion represents the best kind of interplay between theory and experiment. The TD model provided a framework for understanding the function of DA neurons. This framework included the idea of reward prediction error, which raised the question of whether such a prediction can take into account context. Experiments confirmed that the DA neurons do appear to take context into account but in a manner that is inconsistent with the TD model. Thus, the model should be improved.

How the model should be improved will certainly depend on the context. The authors proposed a modified version of the TD model that could account for their results. However, their model was specific to the context they tested, namely how reward probability depends on the recent history of rewards. In principle, “context” is a much broader concept that encompasses any piece of information that can influence the probability of a reward (including, for example, how an estimate of that probability could determine behavior that, in turn, influences whether reward is given or not; Dayan and Balleine, 2002). Can the prediction error encoded by DA neurons take into account any context that the subject can detect? Can models of DA neuron function be expanded to capture the generality of such computations?

As we learn more about the computations performed by DA neurons, we will undoubtedly learn more about how those computations are used by the rest of the brain. The results from this study suggest that these neurons do not compute a restricted form of reward prediction error but rather have access to other information available to the brain that can help to generate and evaluate predictions. Accordingly, these computations are likely to play a variety of roles, not just for reinforcement learning but also for goal-directed behavior, decision making, perception, and other higher functions that could profit from the prediction and evaluation of salient, rewarding events (Montague and Berns, 2002). The impressive confluence of approaches currently being applied to these problems, including machine learning, behavior, economics, and physiology, should help to elucidate the many contexts that are used to perform these computations and the many contexts in which the output of these computations are used.

Joshua I. Gold and Rishi M. Kalwani
Department of Neuroscience
University of Pennsylvania
Philadelphia, Pennsylvania 19104

Selected Reading

- Dayan, P., and Balleine, B.W. (2002). *Neuron* 36, 285–298.
- Fiorillo, C.D., Tobler, P.N., and Schultz, W. (2003). *Science* 299, 1898–1902.
- Montague, P.R., and Berns, G.S. (2002). *Neuron* 36, 265–284.
- Nakahara, H., Itoh, H., Kawagoe, R., Takikawa, Y., and Hikosaka, O. (2004). *Neuron* 41, this issue, 269–280.
- Schultz, W. (2002). *Neuron* 36, 241–263.
- Sutton, R.S., and Barto, A.G. (1981). *Psychol. Rev.* 88, 135–170.
- Sutton, R.S., and Barto, A.G. (1998). *Reinforcement Learning: An Introduction* (Cambridge, MA: MIT Press).
- Wise, R.A. (1996). *Curr. Opin. Neurobiol.* 6, 243–251.
- Wise, R.A. (2002). *Neuron* 36, 229–240.

Language Abilities of Motor Cortex

A new exploration of the cortical network underlying our language abilities by Hauk et al., in this issue of *Neuron*, shows that the process of giving meaning to words differentially activates the motor cortex according to the semantic category of the word.

Understanding the meaning of words that relate to a motor action, such as “dance,” may need more than the well-known language areas of Broca and Wernicke in the left hemisphere of the brain. In this issue of *Neuron*, Hauk et al. (2004) report the surprising discovery that the mere reading of action-related words also activates the motor homunculus—a cortical region of the brain that controls voluntary movements of our different body parts. Remarkably, just the reading of feet-related action

words such as “dance” makes this motor homunculus move its feet.

Language constitutes one of the defining aptitudes of the human species. Almost every aspect of our cognitive life is related to language in the sense that most of our thoughts are linguistically structured (Hauser et al., 2002). The question of which areas of the brain participate in the process of understanding word meaning is among the most intriguing questions in the quest for the physiological explanation of our linguistic abilities.

To gain insight into the brain areas engaged in the comprehension of words, Hauk et al. (2004) designed a reading task in which human subjects were asked to passively read single words presented on a computer screen. The words they carefully chose implied actions performed with different body parts, such as lick, pick, or kick, which are related to movements involving the face, arm, and leg, respectively. At the same time that subjects were reading this type of action words, images of their brains were taken by functional magnetic resonance imaging (fMRI), a noninvasive imaging technique that is used to measure the blood flow changes in the brain that temporally and spatially accompany the activation of neurons.

As can be imagined, several brain structures are active in any behavioral state at any time. So, to obtain an image of the brain regions exclusively involved in a specific cognitive function, the first step is to get a “background” or “basal” image showing all those active areas that are not related to the cognitive function of interest. By means of subtracting this basal image from that obtained when the subject was performing a particular task, the fMRI technique allows the experimenter to locate the brain areas that were active exclusively during the task and that may support the cognitive ability displayed by the subject. By subtracting the image of the brain obtained when the subjects were viewing meaningless hash marks from that obtained when the subjects were reading meaningful words, Hauk et al. (2004) observed that action words, as expected, engaged the well-known classic language areas of the left temporal and frontal cortices. Resonance images also showed activation of the premotor cortex, confirming results obtained previously showing that the cognitive function of identifying and naming objects related to a motor action involves brain areas also engaged in the performance of body movements (Martin et al., 1996).

In addition to the left temporal and frontal cortices, the resonance images revealed a fascinating discovery: the reading of action words also activated the primary motor cortex. Moreover, each of the different groups of words related to actions performed with the face, arms, or legs seemed to engage the motor cortex in a somatotopic fashion (see Figure 1).

To rigorously test the observation of the differential activation of the motor cortex according to the semantic category of the words, the authors superimposed the images of the areas engaged during the reading of the three word groups with images of the motor cortex obtained when the subjects were asked to move their feet, fingers, and tongue. The results of this analysis clearly demonstrated that the motor cortex activated by reading leg- and arm-related words was precisely the same cortex that the subjects used to move their feet and fingers.

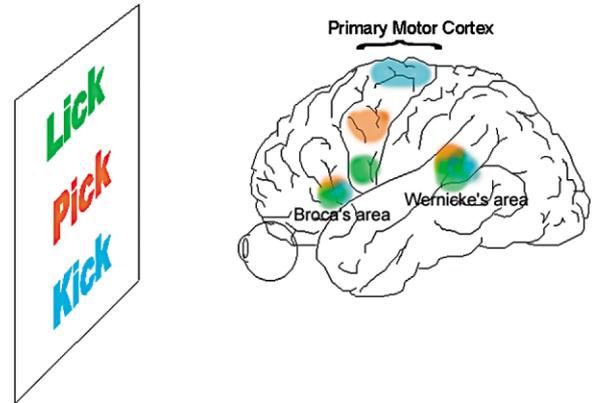


Figure 1. Human Cortical Areas Activated during the Reading of Action Words

The processing of words related to an action performed with a body part, in addition to the classic areas of Wernicke and Broca, selectively activates the primary motor cortex according to the semantic content of the word. For example, words like lick, pick, or kick engage the motor regions related to the control of the tongue, arm, and leg, respectively (blue, red, and green patches).

This revealing discovery means that the same brain area that we use to make a particular body movement is also involved during the process of giving meaning to words that are semantically related to that body part.

The motor and premotor cortices are traditionally thought to participate in the initiation and execution of voluntary movements, storing general motor plans (Graziano et al., 2002), and controlling parameters of motion such as the direction, speed, and force (Georgopoulos, 1999). The involvement of the primary motor area in language processes was unexpected and surprisingly adds a new linguistic dimension to this area. In agreement with a growing body of evidence, this result suggests that the view of the motor areas as the “output stage” of the cortex must be substituted since these cortices are the seat of an extraordinary variety of processing operations (Naito et al., 2002; Salinas and Romo, 1998).

The results of Hauk et al. (2004) suggest that the cortical network supporting language is not localized in single areas but may involve widely distributed areas, differentially activated according to the semantic content of the word. As we have seen, the process of getting the meaning of a word engages the premotor and primary motor areas. An important question, however, is yet to be answered: is the reading task activating the motor areas just because an action and its name commonly cooccur in time (i.e., they are temporally associated), or is this a functional relation in which the motor areas play an active role in comprehension? To explore the issue of a mere statistical relation versus an active relationship between language and motor areas, other techniques than fMRI may be needed. A promising technique that allows some degree of manipulation of the human cortical processing is transcranial magnetic stimulation in which localized, intense but brief magnetic pulses are used to interfere and reversely deactivate a given cortical area (Walsh and Cowey, 2000). By means of transiently deactivating the motor and premotor cortex dur-

ing a linguistic task, the hypothesis of the necessity of these areas in language comprehension could be tested. This experiment could rule out the question of the activation of the primary motor cortex being only accessory and not necessary for the understanding of action words.

The extraordinary results of Hauk et al. (2004) open other intriguing questions such as the following. Which areas participate in the comprehension of words not related to an action or to actions that some of us may have never performed, such as skiing? Do they involve visual imagery and the visual cortex? Would the mere reading of sound-related words activate the auditory cortices? Even more intriguing: what about the neuronal correlates of words related to abstract concepts, such as happiness?

This new vision of the language network, revealing that the human brain areas that are used to perform an action are also needed to comprehend the words related to that action, is consistent with neuronal activity recorded in the macaque brain. Even though these animals lack a developed language, that is, the ability to combine a finite number of words to generate an infinite number of concepts (an ability often called recursion; Hauser et al., 2002), they do seem to possess a basic prerequisite for language: internal representations of actions performed by other individuals. It has been shown that within the premotor cortex there are neurons that are active not only when the monkeys perform an action (Romo et al., 2004) but also when they see other individuals, even the experimenters, performing that action (Rizzolatti and Luppino, 2001). These neurons even activate when the monkeys only listen to sounds produced by the action (such as ripping paper) and are called multimodal mirror neurons (Kohler et al., 2002). These premotor mirror neurons might constitute the neuronal basis that “code and represent abstract concepts” needed in human language and might comprise the neuronal elements recorded in the frontal motor cortices during linguistic tasks (Rizzolatti et al., 2002).

By showing that the comprehension of action-related words engages the motor cortex, Hauk et al. (2004) added a new node to the language cortical network and may also help to transform our view of the motor homunculus. Importantly, these results insinuate to us that the processes of understanding and acting are not terribly different from each other. In this sense, the next time that someone invites you to dance, remember that no matter what your answer is, if you understood the question, your motor cortex already accepted the invitation.

Victor de Lafuente and Ranulfo Romo
Instituto de Fisiología Celular
Universidad Nacional Autónoma de México
04510 México, D.F.
México

Selected Reading

Georgopoulos, A.P. (1999). *News Physiol. Sci.* 14, 64–68.
Graziano, M.S.A., Taylor, C.S.R., Moore, T., and Cooke, D.F. (2002). *Neuron* 36, 349–362.

Hauk, O., Johnsrude, I., and Pulvermüller, F. (2004). *Neuron* 41, this issue, 301–307.

Hauser, M.D., Chomsky, N., and Fitch, W.T. (2002). *Science* 298, 1569–1579.

Kohler, E., Keysers, C., Umiltà, M.A., Fogassi, L., Gallese, V., and Rizzolatti, G. (2002). *Science* 297, 846–848.

Martin, A., Wiggs, C.L., Ungerleider, L.G., and Haxby, J.V. (1996). *Nature* 379, 649–652.

Naito, E., Roland, P.E., and Ehrsson, H.H. (2002). *Neuron* 36, 979–988.

Rizzolatti, G., and Luppino, G. (2001). *Neuron* 31, 889–901.

Rizzolatti, G., Fogassi, L., and Gallese, V. (2002). *Curr. Opin. Neurobiol.* 12, 149–154.

Romo, R., Hernández, A., and Zainos, A. (2004). *Neuron* 41, 165–173.

Salinas, E., and Romo, R. (1998). *J. Neurosci.* 18, 499–511.

Walsh, V., and Cowey, A. (2000). *Nat. Rev. Neurosci.* 1, 73–79.