

Dynamic Remapping

Alexandre Pouget and Terrence J. Sejnowski

Introduction

The term *dynamic remapping* has been used in many different ways, but one of the clearest formulations of this concept comes from the mental rotation studies by Georgopoulos et al. (1989; see also REACHING: CODING IN MOTOR CORTEX). In these experiments, monkeys were trained to move a joystick in the direction of a visual stimulus, or 90° counterclockwise from it. The brightness of the stimulus indicated which movement was required on a particular trial: a dim light corresponding to a 90° movement and a bright light to a direct movement. An analysis of reaction time suggested that, by default, the initial motor command always pointed straight at the target and then continuously rotated if the cue indicated a 90° rotation, an interpretation that was subsequently confirmed by single-unit recordings.

The term *remapping* is also commonly used whenever a sensory input in one modality is transformed to the sensory representation in another modality. The best-known example in primates is the remapping of auditory space, which is head-centered in the early stages of auditory processing, into the retinotopic coordinates used in the superior colliculus (Jay and Sparks, 1987; Stein and Meredith, 1993). This type of remapping, equivalent to a change of coordinates, is closely related to sensorimotor transformations. It does not have to be performed over time but could be accomplished by the neuronal circuitry connecting different representations.

This review is divided into three parts: In the first part, we briefly describe the types of cortical representations typically encountered in dynamic remapping. We then summarize the results from several physiological studies in which it has been possible to characterize the responses of neurons involved in temporal and spatial remappings. Finally, in the third part, we review modeling efforts to account for these processes.

Neural Representation of Vectors

A saccadic eye movement toward an object in space can be represented as a vector, S , whose components s_x and s_y correspond to the horizontal and vertical displacement of the eyes. Any sensory or motor variable can be represented by similar vectors. There are two major ways of representing a vector in a neural population: by a map and by a vectorial representation.

The encoding of saccadic eye movements in the superior colliculus is an example of a map representation. A saccade is specified by the activity of a two-dimensional layer of neurons (see COLLICULAR VISUOMOTOR TRANSFORMATIONS FOR SACCADES). Before a saccade, a bump of activity appears on the map at a location corresponding to the horizontal and vertical displacement of the saccade.

A vectorial code appears to be used for the coding of the direction of hand movement in the primary motor cortex. Georgopoulos et al. (1989) showed that neurons in the primate

motor cortex of a monkey moving its hand fire maximally for a particular direction of the hand movement and respond with a cosine tuning as a function of angle from this best direction. This suggests that each cell encodes the projection of the vector along its preferred direction.

In both cases, the original vector can be recovered from the population activity pattern through a simple transformation. In the map, the center of mass of the activity pattern codes for the direction of movement. For the vectorial representation, an estimate of the original vector, called a population vector, can be recovered by having each unit vote for its best direction by an amount proportional to its activity.

Neurophysiological Correlates of Remapping

Continuous Remappings

Georgopoulos et al. (1989) studied how the population vector varies over time in the mental rotation experiment described in the introduction. They found that, for 90° movements, the vector initially pointed in the target direction and then continuously rotated 90° counterclockwise, at which point the monkey initiated a hand movement (Figure 1A). This is consistent with the interpretation of the reaction time experiments: the monkey had initially planned to move toward the stimulus, and then it updated this command according to stimulus brightness.

Evidence for continuous remappings also has been reported in a double-saccade paradigm. In these experiments, two targets are briefly flashed in succession on the screen, and the monkey makes successive saccades to their remembered locations (Figure 1B). Monkeys can perform this task with great accuracy, demonstrating that they do not simply keep a trace of the retinotopic location of the second target, since, after the first eye movement, this signal no longer corresponds to where the target was in space. Single-unit recordings in the superior colliculus, frontal eye field, and parietal cortex have shown that, before the first saccade, the brain encodes the retinotopic location of the second target. Then, while the first eye movement is executed, this information is updated to represent where the second target would appear on the retina after the first saccade (Mays and Sparks, 1980; Gnadt and Andersen, 1988; Goldberg and Bruce, 1990). In certain cases, this update is predictive, i.e., it starts before the eye movement (Duhamel, Colby, and Goldberg, 1992).

If these remappings were continuous, a bump of activity should sweep through the superior colliculus, going from the retinal location of the second target to its new retinal location after the first saccade. Notice that if a similar mechanism operates on the hill of activity related to the first target, this hill would move toward the center of the map, since the first saccade results in the foveation of the first target. Munoz, Pelisson, and Guitton (1991) have reported evidence for this moving-hill mechanism in single-saccade paradigm, but it is an




Figure 1. A, Rotation of population vector in the primary motor cortex when the brightness of the target (star) indicates a 90° clockwise movement. (Adapted from Georgopoulos et al., 1989.) B, Saccade remapping. The monkey makes a double saccade to the remembered positions of T1 and T2. The post-stimulus-time histograms at the bottom show the responses of two cells with receptive fields RF1 and RF2. The second cell (RF2) responds only after the first eye movement, encoding the new retinal location of T2, even though it is no longer present on the screen.

open question whether a moving hill is also involved in the remapping of the second target.

Although the data from Munoz et al. (1991) and from Georgopoulos et al. (1989), are consistent with continuous remappings, they do not constitute proof of this concept. Hence, the population vector rotation could be the consequence of the simultaneous decay and growth of, respectively, the initial planned hand direction and the final one, without ever truly activating intermediate directions. There are also several unresolved problems with the moving-hill hypothesis of Munoz et al. (1991), as emphasized by Sparks (1993).

One-Shot Sensory Remapping

In the inferior colliculus and primary auditory cortex, neurons have bell-shaped auditory receptive fields in space whose position is fixed with respect to the head. In contrast, in the multi-sensory layer of the superior colliculus, the position of the auditory receptive fields is fixed in retinotopic coordinates, which implies that the auditory map must be combined with eye position (Jay and Sparks, 1987). Therefore, the auditory space is remapped in visual coordinates, presumably for the purpose of allowing auditory targets to be foveated by saccadic eye movements, a function mediated by the superior colliculus.

A similar transformation has been found in the striatum and the premotor cortex, where some of the cells have visual receptive fields in somatosensory coordinates (skin-centered) (Graziano and Gross, 1993). In all cases, these remappings are thought to reflect an intermediate stage of processing in sensorimotor transformations.


These remappings can be considered as a change of coordinates, which correspond to a translation operation (see also SENSOR FUSION). For example, the auditory remapping in the superior colliculus requires the retinal location of the auditory stimulus, R , which can be computed by subtracting its head-centered location, A , from the current eye position, E :

$$R = A - E \quad (1)$$

Remapping Models

The remappings we have described so far fall into two categories: vector rotation with a vectorial code (e.g., mental rotation) and vector translation within a topographic map (e.g., auditory remapping in the superior colliculus). These transformations are very similar, for rotating a vector within a vectorial representation consists of translating a pattern of activity along a circle. Therefore, in both cases, the remapping involves

Figure 2. Remappings as a moving hill of activity in a map. In continuous remapping (left), a recurrent network dynamically moves the hill of activity according to a velocity signal, \dot{E} . In feedforward remapping (right), the hill is moved in one shot by the full amount of the current displacement, E , through an intermediate stage of processing in the hidden layer.



translating a bell-shaped pattern of activity across a map. Most models perform this operation either dynamically through time or in one shot through the hidden layer of a feedforward network (Figure 2).

Dynamical Models

Two kinds of mechanisms have been used in models of continuous remapping: the integration of a velocity signal or the relaxation of a recurrent network.

Integrative model for remapping. In the double-saccade paradigm just described, the retinal coordinates of the second target were updated during the first saccade, a process which might involve moving a hill of activity across the deep layer of the superior colliculus (Munoz et al., 1991). A model by Droulez and Berthoz (1991) shows how this bump of activity could be moved continuously across the map by integrating the eye velocities during the first saccade. Their model is essentially a *forward* model of motion; i.e., given a velocity signal, it generates the corresponding moving image. Interestingly, the equations are very similar to those used for *inverse* models of motion processing. In both cases, the analysis relies on the assumption that the overall gray-level profile in the image is unchanged; only the position of the image features is changed. It is possible to design a recurrent network to implement this constraint (Droulez and Berthoz, 1991). The resulting network moves arbitrary patterns of activity in response to an instantaneous velocity signal.

This model would not only update the coordinates of second target, but would also move the hill of activity corresponding to the first target toward the center of the map, consistent with Munoz et al. (1991). Finally, the model does not require the true eye velocity to work. An approximation of eye velocity,

obtained from the eye position modulated neurons found in the parietal cortex (see GAZE CODING IN THE POSTERIOR PARIETAL CORTX), would be sufficient (Dominey and Arbib, 1992).

Recurrent networks. Mental rotation of a population vector can be reproduced by training a neural network to follow a circular trajectory over time. In this case, the population vector rotates as a consequence of the network dynamics in the absence of any input signals. This approach has been used by Lukashin and Georgopoulos (1994) to model the generation of hand trajectory, but when the trajectory is a circle, mental rotation and a circular hand trajectory are equivalent. Although the model generates a rotating vector, additional mechanisms must be specified to stop the rotation.

Single-Shot Models

Feedforward models have been used for vectorial as well as map representations. They are used whenever the amplitude of the shift is available to the brain beforehand, as for auditory remapping in the superior colliculus in which the shift is directly proportional to the current eye position (Equation 1). In contrast, in the mental rotation case, the amplitude of the shift is specified by an external stimulus.

Shifter models. As demonstrated by Touretzky, Redish, and Wan (1993), rotation within a vectorial representation can be performed by using a shifter circuit (for more details on shifter circuits, see also ROUTING NETWORKS IN VISUAL CORTEX). Their architecture uses N independent circuits, each implementing a rotation through a particular angle. This mechanism is limited in resolution since it rotates only in multiples of $360/N$ degrees. Whether such shifter circuits actually exist in the brain remains to be demonstrated.

Feedforward network models. There are multiple examples of three-layer networks, or variations thereof, which are trained or hand-crafted to perform sensory remapping. Since these remappings perform vector addition, it might appear unnecessary to deploy a fully nonlinear network for such a task. However, with a map representation, vector addition requires moving a hill of activity in a map as illustrated in Figure 2, an operation that is highly nonlinear.

Special-purpose nonlinear circuits can be designed to perform this operation (Groh and Sparks, 1992), but more biologically realistic solutions have been found with networks of sigmoidal units trained with backpropagation. Hence, the model of Zipser and Andersen (see GAZE CODING IN THE POSTERIOR PARIETAL CORTEX), which was trained to compute a head-centered map from a retinotopic input, uses hidden units with retinotopic receptive fields modulated by eye position, as in parietal neurons (see also Krommenhoek et al., 1993).

However, backpropagation networks are, in general, quite difficult to analyze, providing realistic models but not much insight into the algorithm used by the network. Pouget and Sejnowski (1995) have recently explored a way to analyze such networks using the theory of basis functions.

Basis functions. The process of moving a hill of activity in a single shot can be better understood when considered within the larger framework of nonlinear function approximation. Consider, for example, the feedforward network shown in Figure 2 applied to a remapping from retinotopic, r_x , to head-centered coordinates, a_x . It can be shown that the responses of the output units are nonlinear in the input variables, namely, the retinal position, r_x , and eye position, e_x .

Therefore, the actual goal of the network is to find an appropriate intermediate representation to approximate this output function. One possibility is to use basis functions of r_x and e_x in the hidden layer (Pouget and Sejnowski, 1995).

Perhaps the best-known set of basis functions is the set of cosine and sine functions used in the Fourier transform. Another example is the set of Gaussian or radially symmetric functions with local support (see RADIAL BASIS FUNCTION NETWORKS). A good model of the response of parietal neurons, which are believed to be involved in remapping, is a set of Gaussian functions of retinal position multiplied by sigmoids of eye position. The resulting response function is very similar to that of gain-modulated neurons in the posterior parietal cortex (see GAZE CODING IN THE POSTERIOR PARIETAL CORTEX).

Conclusions

Remapping can be continuous and dynamic or a single shot through several layers of neurons. In both cases, the problem amounts to moving a hill of activity in neuronal maps. Whether some models are better than others is often difficult to establish because the neurophysiological data available are still relatively sparse. Modelers might be able to constrain their models

further by considering deficits that have been documented after lesions in humans. These data not only provide valuable insights into the nature of remappings, but they also might help bridge the gap between behavior and single-cell responses.

Road Maps: Primate Motor Control; Vision

Related Reading: Corollary Discharge in Visuomotor Coordination

References

- Dominey, P. F., and Arbib, M. A., 1992, A cortico-subcortical model for the generation of spatially accurate sequential saccades, *Cereb. Cortex*, 2:153-175.
- Droulez, J., and Berthoz, A., 1991, A neural model of sensorimotor maps with predictive short-term memory properties, *Proc. Natl. Acad. Sci.*, 88:9653-9657. ♦
- Duhamel, J. R., Colby, C. L., and Goldberg, M. E., 1992, The updating of the representation of visual space in parietal cortex by intended eye movements, *Science*, 255:90-92.
- Georgopoulos, A. P., Lurito, J. T., Petrides, M., Schwartz, A. B., and Massey, J. T., 1989, Mental rotation of the neuronal population vector, *Science*, 243:234-236. ♦
- Gnadt, J. W., and Andersen, R. A., 1988, Memory related motor planning activity in posterior parietal cortex of macaque, *Exp. Brain Res.*, 70(1):216-220.
- Goldberg, M. E., and Bruce, C. J., 1990, Primate frontal eye fields: III. Maintenance of a spatially accurate saccade signal, *J. Neurophysiol.*, 64(2):489-508.
- Graziano, M. S., and Gross, C. G., 1993, A bimodal map of space: Somatosensory receptive fields in the macaque putamen with corresponding visual receptive fields, *Exp. Brain Res.*, 97(1):96-109.
- Groh, J. M., and Sparks, D. L., 1992, Two models for transforming auditory signals from head-centered to eye-centered coordinates, *Biol. Cybern.*, 67(4):291-302. ♦
- Jay, M. F., and Sparks, D. L., 1987, Sensorimotor integration in the primate superior colliculus: I. Motor convergence, *J. Neurophysiol.*, 57:22-34.
- Krommenhoek, K. P., Van Opstal, A. J., Gielen, C. C. A. M., and Van Gisbergen, J. A. M., 1993, Remapping of neural activity in the motor colliculus: A neural network study, *Vision Res.*, 33(9):1287-1298.
- Lukashin, A. V., and Georgopoulos, A. P., 1994, A neural network for coding trajectories by time series of neuronal population vectors, *Neural Computat.*, 6(1):19-28.
- Mays, L. E., and Sparks, D. L., 1980, Dissociation of visual and saccade-related responses in superior colliculus neurons, *J. Neurophysiol.*, 43(1):207-232. ♦
- Munoz, D. P., Pelisson, D., and Guitton, D., 1991, Movement of neural activity on the superior colliculus motor map during gaze shifts, *Science*, 251:358-360.
- Pouget, A., and Sejnowski, T. J., 1995, Spatial representation and basis functions (in preparation). ♦
- Sparks, D. L., 1993, Are gaze shifts controlled by a "moving hill" of activity in the superior colliculus? *Trends Neurosci.*, 16(6):214-218.
- Stein, B., and Meredith, M. A., 1993, *The Merging of the Senses*, Cambridge, MA: MIT Press.
- Touretzky, D. S., Redish, A. D., and Wan, H. S., 1993, Neural representation of space using sinusoidal arrays, *Neural Computat.*, 5:869-884.