

- synapses in rat medial prefrontal cortex, *J. Neurophysiol.*, 83:3031–3041.
- Mainen, Z. F., and Sejnowski, T. J., 1995, Reliability of spike timing in neocortical neurons, *Science*, 268:1503–1506.
- Senn, W., Markram, H., and Tsodyks, M., 2001, An algorithm for modifying neurotransmitter release probability based on pre- and postsynaptic spike timing, *Neural Computation*, 13:35–67.
- Shadlen, M. N., and Movshon, J. A., 1999, Synchrony unbound: A critical evaluation of the temporal binding hypothesis, *Neuron*, 24:67–77.
- Singer, W., 1999, Neuronal synchrony: A versatile code for the definition of relations? *Neuron*, 24:49–65.
- von der Malsburg, C., 1981, *The Correlation Theory of Brain Function*, MPI Biophysical Chemistry, Internal Report 81–2, reprinted in *Models of Neural Networks II* (E. Domany, J. L. van Hemmen, and K. Schulten, Eds.), Berlin: Springer-Verlag, 1994, chap. 2, pp. 95–119. ♦
- von der Malsburg, C., 1985, Nervous structures with dynamical links, *Ber. Bunsenges. Phys. Chem.*, 89:703–710.
- von der Malsburg, C., 1986, Am I thinking assemblies? in *Proceedings of the Trieste Meeting on Brain Theory, October 1984* (G. Palm and A. Aertsen, Eds.), Berlin: Springer-Verlag, pp. 161–176.
- von der Malsburg, C., 1999, The what and why of binding: The modeler's perspective, *Neuron*, 24:95–104.
- von der Malsburg, C., and Bienenstock, E., 1987, A neural network for the retrieval of superimposed connection patterns, *Europhys. Lett.*, 3:1243–1249.
- von der Malsburg, C., and Schneider, W., 1986, A neural cocktail-party processor, *Biol. Cybern.*, 54:29–40.
- Wiskott, I., and von der Malsburg, C., 1996, Face recognition by dynamic link matching, in *Lateral Interactions in the Cortex: Structure and Function* (J. Sirosh, R. Miikkulainen, and Y. Choe, Eds.), electronic book, available: <http://www.cs.utexas.edu/users/nn/web-pubs/htmlbook96/>.
- Zhu, J., and von der Malsburg, C., 2001, Synapto-synaptic interactions speed up dynamic link matching, presented at the Computational Neuroscience Meeting (CNS*01), San Francisco, CA, June 30–July 5, 2001.

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 MOTOR CORTEX: CODING AND DECODING OF DIRECTIONAL OPERATIONS). 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 corresponded 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 a sensory representation in another modality. The best-known example in primates is the remapping of auditory space, from head-centered in the early stages of auditory processing to the retinotopic coordinates used in the superior colliculus (Jay and Sparks, 1987). 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 where 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 a similar vector. There are two major ways of representing a vector in a neural population—by a topographic map and by a nontopographic vectorial representation.

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

Another example of a vectorial code is the code for the direction of hand movements in the primate motor cortex. Neurons in the primary motor cortex respond maximally for a particular direction of hand movement with a cosine tuning curve around this preferred direction (Georgopoulos et al., 1989). This suggests that each cell encodes the projection of the vector along its preferred direction. [Todorov (2000) questions this interpretation, but the precise identity of the vector being encoded in motor cortex is not critical to the issue of remapping.]

In both cases, the original vector can be recovered from the population activity pattern using statistical estimators. Various examples of such estimators are described in POPULATION CODES.


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 movement 90° counterclockwise from the target, the vector encoded in M1 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 updated this command according to the task requirement.

Similar continuous remapping occurs in the postsubiculum of the rat, one of the cortical structures involved in navigation of space. Neurons in the postsubiculum provide an internal compass that encodes the direction of the head with respect to remembered visual landmarks. The neurons have bell-shaped tuning curves around their best direction, similar to the code for hand direction in the primary motor cortex. Electrophysiological recordings have revealed that this vector is continuously updated as the head of the

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 (S1 and S2) to the remembered positions of T1 and T2. C, Post-stimulus-time histograms showing the responses of two cells with receptive fields RF1 and RF2 illustrated in Figure 1B. 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.



rat moves in space, even in complete darkness, suggesting that vestibular inputs are used for this updating (see RODENT HEAD DIRECTION SYSTEM).

Another example of continuous remappings has been reported in a double saccade task. 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 the brain encodes the retinotopic location of the second target before the first saccade occurs. 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 (Figure 1C; Mays and Sparks, 1980). In certain cases, this update is predictive; i.e., it starts prior to the eye movement (Duhamel, Colby, and Goldberg, 1992).

Graziano, Hu, and Gross (1997) have reported that the same mechanism appears to be at work in the premotor cortex. Bimodal, visuotactile neurons with receptive fields on the face remap the position of remembered visual stimuli after head movements. It is therefore becoming increasingly clear that continuous remappings

are widespread throughout the brain and play a critical role in sensorimotor transformations.

Although all these examples clearly involve vector remappings, it is not entirely clear that the remappings are continuous. Hence, in the Georgopoulos et al. (1989) experiment, the population vector rotation could be a consequence of the simultaneous decay and growth of the initial planned hand direction and the final one, respectively, without ever activating intermediate directions. This is an example of one-shot remapping considered in the next section. Moreover, it is often difficult to determine whether a remapping in one particular area is computed in that area or is simply the reflection of a remapping in an upstream area.

One-Shot Sensory Remapping

In the inferior colliculus and primary auditory cortex, neurons have bell-shaped auditory receptive fields in space whose positions are fixed with respect to the head. In contrast, in the multisensory layer of the superior colliculus, the positions of the auditory receptive fields are 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 et al., 1997). 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. For example, the auditory remapping in the superior colliculus requires the retinal location of the auditory stimulus, R , which, to a first approximation, 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 similar, since rotating a vector within a vectorial representation consists of translating a pattern of activity around a circle. Therefore, in both cases the remapping involves 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 described above, the retinal coordinates of the second target were

updated during the first saccade, a process that might involve moving a hill of activity within the parietal cortex. 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 (Figure 1A). Their model is essentially a *forward* model of motion: Given a velocity signal, it generates the corresponding moving image. Interestingly, the equations are similar to those used for *inverse* models of motion processing. In both cases, the analysis relies on the assumption that the temporal derivative of a moving image is zero. In other words, the overall gray level profile in the image is unchanged; only the positions of the image features change. It is possible to design a recurrent network to implement this constraint (Droulez and Berthoz, 1991), and the resulting network moves arbitrary patterns of activity in response to an instantaneous velocity signal.

Several variations of this idea have been developed. Dominey and Arbib have shown that an approximation of eye velocity, obtained from the eye position modulated neurons found in the parietal cortex is sufficient for this architecture to work (Dominey and Arbib, 1992). Their simulations show patterns of activation very similar to the ones shown in Figure 1B in the part of their model corresponding to the parietal cortex, FEF, and superior colliculus. Zhang (1996) has used line attractor networks to model head direction cells in the postsubiculum of the rat. In this model, the hill is moved by using the velocity signal—in this case a head velocity signal—to temporarily modify the efficacy of the lateral connections.

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 Georgopou-




Figure 2. In a map representation, remappings involve moving hills of activity. These hills can be moved continuously in a recurrent network (A), or in one shot in a feedforward network (B). A, The recurrent network dynamically moves the hill of activity according to a velocity signal, E . As described in the text, there are several ways to achieve this result. Droulez and Berthoz (1991) integrate the eye velocity signals through the lateral connections while Zhang (1996) uses the eye velocity signals to temporarily bias the lateral connections. B, In feedforward remapping, the hill is moved in one shot by the full amount of the current displacement, E , via an intermediate stage of processing in the hidden layer. The weights can be adjusted with a learning algorithm such as backpropagation. Alternatively, one can use basis function units in the hidden layer and train the weights to the output units with a simple learning algorithm such as the delta rule.

os (1994) to model the generation of hand trajectories, 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, such as auditory remapping in the superior colliculus in which the shift is directly proportional to the current eye position (Equation 1). In contrast, for mental rotation, 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 ROUTING NETWORKS IN VISUAL CORTEX in the First Edition). Their architecture uses N independent circuits, each implementing a rotation through a particular angle. This mechanism is limited in resolution since it rotates only by multiples of $360/N$ degrees. Whether such shifter circuits actually exist in the brain remains to be demonstrated.

Feedforward network models. There are many examples of three-layer networks, and variations thereof, that have been trained or handcrafted to perform sensory remappings. 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 2B, 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 in the First Edition), 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 generally quite difficult to analyze, providing realistic models but little insight into the algorithm used by the network. Pouget and Sejnowski (2001) have 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. For example, consider the feedforward network shown in Figure 2B, applied to a remapping from retinotopic, R_x , to head-centered coordinates, A_x . Because of the map format used in the output later, 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, 2001; Salinas and Abbot, 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 sigmoid functions 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 in the First Edition, and Pouget and Snyder (2000) for a review].

Conclusions

Remappings 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 simply because the neurophysiological data available are relatively sparse. Models can be further constrained by considering deficits that accompany localized lesions in humans (see Pouget and Sejnowski, 2001). These data not only provide valuable insights into the nature of remappings but also might help bridge the gap between behavior and single-cell responses.

Road Map: Vision

Related Reading: Collicular Visuomotor Transformations for Gaze Control; Motion Perception: Elementary Mechanisms; Pursuit Eye Movements; Visual Attention; Visual Scene Perception

References

- Dominey, P., and Arbib, M., 1992, A cortico-subcortical model for the generation of spatially accurate sequential saccades, *Cerebral Cortex*, 2:153–175. ◆
- Droulez, J., and Berthoz, A., 1991, A neural model of sensoritopic maps with predictive short-term memory properties, *Proc. Natl. Acad. Sci. USA*, 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(5040):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.
- Graziano, M., Hu, X., and Gross, C., 1997, Coding the locations of objects in the dark, *Science*, 277:239–241.
- Groh, J., and Sparks, D., 1992, Two models for transforming auditory signals from head-centered to eye-centered coordinates, *Biol. Cybernetics*, 67:291–302.
- 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., and Van Gisbergen, J. A. M., 1993, Remapping of neural activity in the motor colliculus: A neural network study, *Vision Res.*, 33:1287–1298.
- Mays, L. E., and Sparks, D. L., 1980, Dissociation of visual and saccade-related responses in superior colliculus neurons, *J. Neurophysiol.*, 43:207–232.
- Pouget, A., and Sejnowski, T. J., 2001, Simulating a lesion in a basis function model of spatial representations: Comparison with hemineglect, *Psychol. Rev.*, 108:653–673. ◆
- Pouget, A., and Snyder, L., 2000, Computational approaches to sensorimotor transformations, *Nature Neurosci.*, 3:1192–1198.
- Salinas, E., and Abbot, L., 1995, Transfer of coded information from sensory to motor networks, *J. Neurosci.*, 15:6461–6474. ◆
- Todorov, E., 2000, Direct cortical control of muscle activation in voluntary arm movements: A model, *Nature Neurosci.*, 3:391–398.
- Touretzky, D., Redish, A., and Wan, H., 1993, Neural representation of space using sinusoidal arrays, *Neural Computation*, 5:869–884.
- Zhang, K., 1996, Representation of spatial orientation by the intrinsic dynamics of the head-direction cell ensemble: A theory, *J. Neurosci.*, 16:2112–2126. ◆