tion is achieved anew with each fixation" (Petrov 1992; Petrov & Zenkin 1976b), but unfortunately we find ourselves at a starting position advanced not so far from beloved Von Helmholtz and, again, we have much work to do to explain visual stability across saccades.

Is perception isomorphic with neural activity?

Alexandre Pouget and Terrence J. Sejnowski

alex@helmholtz.sdsc.edu; Howard Hughes Medical Institute Research Laboratories, The Salk Institute, La Jolla, CA 92037

The theory of visual stability proposed by Bridgeman et al. requires the existence of retinotopic maps in which neurons represent the spatiotopic position of objects *through their activity*. The authors reason that if the code for position is fully specified by the neural activity, it becomes independent of the position of neurons within the map (sect. 4.1, para. 3).

These two aspects of Bridgeman et al.'s theory, namely, the retinotopy of the maps and the encoding of position through the activity regardless of retinotopic position, appear difficult to reconcile. Neurons within retinotopic maps in early visual cortex are known to have small receptive fields. Accordingly, a single cell cannot encode all possible retinal positions through its firing rate. If the head and trunk are fixed, the position of an object in space is the vector sum of its retinal position and eye position. If a cell cannot encode all possible retinal positions, it is difficult to imagine how it could do so for spatiotopic positions. This in turn entails that a neuronal population at a given location on the map cannot possibly encode an arbitrary head-centered position through its activity.

The problem stems from the fact that in computing spatiotopic position in retinotopic maps, there is an intimate relation between what Bridgeman et al. call *the mode of representing* and *the position of representing* (sect. 4, para. 1). One relevant variable, retinal position, is itself represented by the map. As emphasized in the target article, this is not a problem for other visual attributes such as color or motion (sect. 4, para. 5–6). Extremely large receptive fields would certainly solve the problem (Zipser & Andersen 1988), but neurons in V1 or V2 do not have such large receptive fields.

Nevertheless, it is still possible that the early visual cortical areas may support, partly or fully, a representation of the visual field that is suitable for determining visual stability during eye movements. We have presented a model for one way this could be done in a recent paper (Pouget et al. 1993), inspired by the reports from several laboratories that eye position modulates visually driven activity in the lateral geniculate nucleus (Lal & Friedlander 1989; Lehmkuhle & Baro 1991) and area V1 and V3a of visual cortex (Galleti & Battaglini 1989; Trotter et al. 1992; Weyand & Malpeli 1993). We have shown that such neurons could support a distributed representation of objects, including their head-centered spatial locations; the location is encoded by the modulation of the visual responses of neurons by eye position. We called this representation a retinospatiotopic map.

The code for spatiotopic position in a retinospatiotopic map depends on both the pattern of activity and the positions of neurons in the map. The eye position is encoded through the neural activity and the retinal position is provided by the place code within the map. Although the codes for both dimensions are different, the network model demonstrates that it is possible to recover the head-centered position of an object in the output layer. The output layer is not essential for the spatial representation. There does not have to be such an invariant representation of spatial location anywhere in the brain and, indeed, none has been reported. We conclude only that spatial information is relatively easy to extract from the populations of neurons, as well

as functions of spatial location such as those that might be needed for guiding hand movements.

This raises a critical difference. We suspect that Bridgeman et al. would not accept retinospatiotopic maps as an account for visual stability. What would constitute an account of visual stability is never explicitly stated in their paper, although from the discussion in section 4.1 we infer that a visual area is responsible for visual stability if, and only if, patterns of neural activity are isomorphic with the spatiotopic position of objects; or, equivalently, if and only if positions in space are in a one-to-one relationship with patterns of neuronal activity. Whether the patterns move to different locations in the brain as a result of eye movements is irrelevant for them, as long as the pattern stays the same. A retinospatiotopic map fails because it does not encode head-centered position solely through an activity pattern but through a combination of activity and topography. The isomorphism criterion may be too strict, since even brain regions such as parietal area 7a and the lateral intraparietal area (LIP), which are believed to represent the position of objects in spatiotopic coordinates, would not qualify according to this strict requirement. Nonetheless, area 7a appears to contain a distributed representation of objects in body-centered coordinates (Andersen & Zipser 1988; Goodman & Andersen 1990; Zipser & Andersen 1988).

It is possible that there is no invariant representation of spatial location anywhere in the brain. Would this make any computational problem that the brain must solve insoluble as a consequence? On the contrary, any invariant function, such as grasp of an object in space, could be computed from a distributed representation. Objects are grasped, brain activity comes and goes, and no invariant representation is to be found, anywhere, except perhaps at the level of the motoneurons. The strong isomorphism hypothesis is a conceptually simple one that would be convenient if true, but by no means necessary. We must begin to imagine how it could be that more than one pattern of activity can represent the same percept. Models provide a way to embody this otherwise difficult to accept possibility.

This leaves open the question of what it is that constitutes the feeling that the world is stable, which may have quite a different answer from the question of how the stable world is represented. Note that stability has a much lower dimensionality than the world itself – in the simplest case we need only a single scalar neuron that monitors stability. Such a stability monitor would need to receive converging inputs from many part of the brain and would act as a comparator for successive brain states. The brainstem would be a more likely place to find such neurons than the cortex. Such a system could be useful for balance and might even contribute to the control of posture. It could not, however, be used to coordinate transformations between sensory modalities, nor would it help in guiding actions. There is no reason different representations of the stable world might not be found in different parts of the brain for different purposes.

Neuronal death of the cancellation theory?

Claude Prablanc

Inpeu94@frsun12.bitnet; INSERM Unité 94, Vision et Motricité, 69500 Bron, France

The question of how the brain can construct a stable representation of the external world despite eye movements is a very old one. If there have been some wrong statements of problems (such as the inverted retinal image), other statements are less naive and have led to analytic solutions possibly adopted by the brain to counteract the spurious effects of eye movements. Following the MacKay (1973) objections to the analytic view of perceptual stability, Bridgeman et al. claim that the idea that signals canceling the effects of saccadic eye movements are