

they are at (even relative to one another) may or may not be represented. The same is true of the location of objects on the proximal stimulus (e.g., on the retina) or further up in the nervous system, such as patterns of activity on the retinotopically organized fibers leading from the eye, or in the primary visual cortex, which is largely retinotopically mapped. Since these locations are past the sensors, are they necessarily representations? If so, what is the essential difference between the way that distance in the world affects perception and the way that the corresponding distance on a neighborhood-preserving (i.e., homeomorphic) anatomical mapping affects perception (for ease of reference I will refer to the results of such mappings as “neural layouts” or NLs)? We can say that such neural layouts *register* (rather than represent) spatial properties. They help to illustrate the general theme that there are many types of representations, ranging from conceptual, through subpersonal, to informational states that are better referred to as registrations rather than representations. In the next subsection I will examine neural layouts to see if they warrant the conclusion that spatial properties are always represented in NLs since locations appear to be roughly preserved on a maplike surface.

(3) Are neural layouts always representations? Intuitively it seems that a neural layout (a layout of activity in the cortex that is a homeomorphic mapping of some other spatial domain, such as shown in figure 4.4) carries information about location in a special way that makes it a maplike representation (I will have more to say about maplike representations and their role in navigation in section 5.4.2). The intuition is that any projection of spatial information onto a neural layout (NL) is automatically a representation since it preserves spatial locations (at least to a first approximation). This intuition derives from the fact that such an NL resembles a canonical map or picture and could (if spread out) be used by a person to navigate or to recognize patterns. However, the layout need not be used in this way.

Whether we call any retinal or other neural layout a *representation* is partly a question of terminology, and NLs usually do carry information about something in the world to which they are causally connected. As mentioned above I prefer to call that type of mapping a *registration* of information—spatial properties are registered rather than represented in NLs. What does matter is not the terminology, but the distinctions we need to make with respect to the role NLs play in explanations. If we use the term “representation” to refer to any information-bearing state, then we will still need to distinguish another, stronger sense of representation. The main distinction we still need is that between states whose representational *con-*

tent plays a role in explanations and those in which the content (if any) does not play any role. If we gain no explanatory advantage by specifying *what* an NL represents, then nothing is gained by treating the NL as a representation. The fact that the NL looks like a map—even if places on the NL correspond to places in the world—is still not enough for it to be a representation in the strong sense.

There are several specific requirements that should be met for something to count as a spatial representation in the strong sense. We need to show not only that locations, distances, and directions in the NL correspond to the same properties in the world but also that they determine the organism's behavior vis-à-vis those represented places. In other words, we need to show that these properties of the NL function to represent properties of the world for the organism. One indicator that they function in this way is if at least some generalizations concerning behavior require appeal to the represented domain as opposed to the pattern of the NL itself. Some principles governing NLs might well be captured solely in terms of properties of the NL with no regard to what they may represent. The principles for forming clusters of features and most Gestalt grouping principles may well be of this sort. These principles (at least as understood by people like Kohler and Wertheimer) are expressed over properties of the proximal stimulus or over neural fields in the brain,⁵ but not over locations, distances, and directions in the world.

One way to see this is to reflect on the fact that unless the function of the NL is to represent spatial properties for the organism, it would not be possible for the NL to *misrepresent* something. The possibility of misrepresenting is a signature property of representations—a retinal pattern cannot misrepresent the visual world since optics does not make “mistakes.” Similarly, it is meaningless to ask of an NL in which frame of reference it represents an

5. Processes operating over NLs typically respond to spatially local properties of the NL—they operate over “local support.” The principles of operation of such processes are *prima facie* expressible over nonconceptual neural properties. Recently there has been an increase of interest in applying dynamic systems theory to modeling the mind. Since such theories are generally not representational (and not computational in the sense discussed in Pylyshyn 1984) there is little chance that they will explain cognitive processes. But they may find application in the sort of nonrepresentational processes that transform NLs or registrations, derive Gestalt clusters, solve the correspondence problem in certain cases, and even carry out tracking (examples are found in Koch and Ullman 1985; Pylyshyn 2003, appendix 5A). Thus theories that postulate spatial registrations may be appropriate for the sort of neural field processes envisioned by Wolfgang Kohler (1947).

object's location, since by itself it does not represent an object as located anywhere. But in the strong sense of representation, where the NL functions to direct movements or to identify objects, it *does* matter how its spatial relations are represented. In that case an NL may represent some locations with respect to a head-centered frame of reference, or as being to the left of another location, or as being more than an arm's length away; and for purposes of determining actions *it matters how the location is represented* (or what it is represented *as*). Without this strong sense of representation, with only a direct object-to-NL mapping, there is no possibility of misrepresentation, and thus it is misleading to call the NL a representation or a map.⁶

It's important to keep in mind that this discussion is about *explaining* regularities in vision and behavior. So the answer to the question at hand—whether an NL is a spatial representation—is that it depends on whether one must refer to the geometrical properties of the represented world in providing explanations. For example, do the principles (such as principles of clustering or of correspondence) that have to be cited refer to properties of the NL or properties in the world? Suppose, for the sake of argument, that the clustering algorithm applies only over distances on the neural lay-

6. I am leaving out a lot here. What makes a terrestrial map able to misrepresent is that this sort of map typically is constructed with the intention that certain of its features correspond to certain features of the relevant terrain, and the map has to be interpreted with these intentions in mind. Thus there is ample room for the intended correspondence to fail and for the map thus to misrepresent. These degrees of freedom are absent in the case of NLs unless we assume that the map is interpreted by some process that allows a possibility of misinterpretation. Sometimes it is tempting to assume an interpreter, and at other times it is tempting to assume a design purpose for the NL—and sometimes it makes sense to talk of a “map” even though there may be no NL, as in the case of insect navigation (see section 5.4.2). Talk about the design purpose (what the NL is *for*) is sometimes helpful, even though strictly speaking there is no agent who designed the representation-using system, because it ties together a variety of otherwise unconnected properties of various mechanisms. In fact our understanding of “natural constraints” rests on assumptions about the purpose of some of the mechanisms, and Marr (1982) motivated his analysis by asking what various visual mechanisms were *for*. Dretske (1981) suggests another way in which an information-carrying state might misrepresent, a way based on learning: If the system has been exposed to pairs of properties and internal states, it could learn which features of the environment the states represent and thus could be in a position to misrepresent those features. These are all questions that I will not get into, beyond arguing that there is more to being a map than homeomorphism.

out, which, in turn, is a homeomorphic transformation of activity on the retina.⁷ In that case nothing is gained by saying that these distances represent properties in the world, since by hypothesis the distance on the NL is all that is relevant to explanations involving distances and those are indifferent between whether it represents a visual angle, a 2-D distance on the retina, a 2-D distance far away from the observer, or a distance in 3-D oriented at the appropriate angle from the viewer to project onto the line on the NL. Therefore, it is not a representation in the strong sense; it does not represent the property *as* something in the world, notwithstanding that, if spread out on a flat surface, the pattern of activity looks like a map. But since it carries some information about spatial locations we say that it *registers* spatial properties.

(4) *When should we postulate representations?* The purpose of postulating representations is to provide explanations and to capture generalizations that would not be captured without reference to the contents of such representations—to what they represent. But sometimes (as in the hypothetical NL discussed above) the function of information-carrying states can be fully described without reference to contents. It could be that principles such as, say, those involved in clustering or apparent motion can be fully explained without reference to any representational content of the states involved. In discussing the way information might be carried by an NL, I noted earlier that an explanation might sometimes be stated in terms

7. These examples are for purposes of illustration; I am not prejudging the empirical question of whether the principles of clustering or of pairing features to solve the correspondence problem apply only to proximity on the NL. If they apply to distal properties then the present argument would not work—but then again neither could we claim that the NL is the basis for the clustering or correspondence solution, since we know at least that V1 (or any other NL) is prior to processes that provide 3-D information (prior to the constancies). There have been conflicting claims over whether 3-D properties are relevant to apparent motion; some investigators maintain that 3-D distance is relevant (Attneave and Block 1973; Wright, Dawson, and Pylyshyn 1987) and some that it is not (Ullman 1979). Recent years have seen many reports of 3-D properties being relevant to what seem like early processes, such as popouts in search (Enns and Rensink 1990; Rensink and Enns 1995) or even multiple object tracking, where it seems that speed in the distal world, rather than on the retina, determines the performance in MOT (Enns and Franconeri 2006; Liu, Austen, Booth, Fisher, et al. 2005). These suggest that such processes are postconstancy or postdepth analysis and therefore do not involve (only) places on the NL (in V1). But this is an empirical question that requires further research.