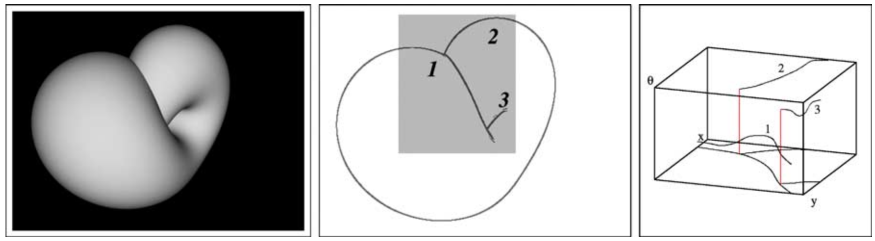

Fig. 15. An illustration of the lift into position (x, y), orientation (θ) space. Three contour fragments from the tangent map in Fig. 5(center) are highlighted. Notice how the discontinuity in orientation at the 1–2 T-junction is separated, highlighting multiple orientations at the same position, the natural columnar representation for orientation discontinuities.

only if they properly exist within the component receptive fields. The large "surround" operator component in the curvature response must arise from a long curve for it to be meaningful; not from an unrelated series of fragments. Furthermore, Dobbins et al. [11] have found that even-symmetry (large) receptive fields provide an estimate of the magnitude of the curvature, while odd-symmetry surround operators provide information about the sign of curvature.

4.3. The position-orientation representation

We now take a step back from the circuitry, to illustrate how this geometry is represented in the columnar machine sketched earlier (Fig. 6). As in that figure, the different orientation possibilities can be viewed as a structure "on top of" the image, with retinotopic (x, y) coordinates extended into a third dimension ("height") corresponding to orientation. This orientation axis is different from the x and y axes, because orientation wraps around 2π with the circle S^1 being its domain. Thus this (position, orientation) space is not modeled as (x, y, z), but as (x, y, θ), where θ is the tangent angle. A point in this space is a point in $R^2 \times S^1$.

It is instructive to consider how different curves in the (x, y) plane lift into $R^2 \times S^1$, to understand some of its advantages for the cortical columnar machine. Not much happens for a straight line in the plane, which lifts to a "horizontal" straight line in $R^2 \times S^1$ at a "height" dependant only on the angle θ. A smooth, closed curve in the plane lifts into a smooth, closed curve in $R^2 \times S^1$. A real difference arises when we consider curves that are continuous but with corners. Such events are important because they could signal a monocular occlusion cue and as we mentioned ealier, all these discontinuities in orientation lift into broken curves in $R^2 \times S^1$ (Fig. 15). In the following we denote it as the (x, y, θ) space.

5. The geometry of texture flows

With this understanding of the inference of tangent maps for individual curves, we move to patterns of multiple curves. Examples include pinstriped material and artist's etchings, animal coats and zebra's stripes. For such patterns orientation is distributed over two-dimensional regions, and as the requirement for perfect continuations is relaxed we obtain texture flows.

The importance of locally (almost) parallel structure has been observed psychologically [14,24,44]. In particular, the human visual system has the tendency to organize and group parallel structure into coherent units. Examples include Kanizsa's "social conformity of a line" [15,24], and Glass's random dot moire patterns [14,41]. In all cases, this organization drastically effects the interpretation of a scene both in 2D and 3D (e.g., [44]).

Informally, texture flows are defined by their orientation content—a dense visual percept characterized by local parallelism and slowly varying dominant local orientation. However, texture flows are not simply a matter of completing a number of almost "parallel" curves, but rather involve dense interpolation of a field of orientations (Fig. 16). This should be compared with the gap-completion property for curves (e.g., [18,25,47]).

Represented as a scalar orientation function, the local behavior of the flow is governed to first-order by its gradient. Unfortunately, this quantity does not capture all of the relevant geometry. As with curves, psychophysical evidence suggests that texture flow perception is linked to curvature [17,30,50,51].

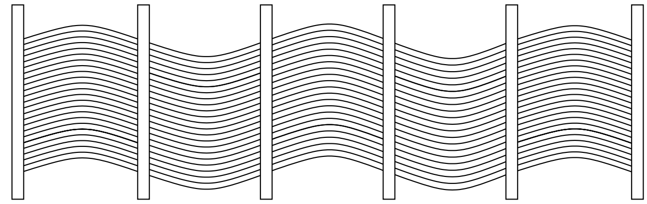

Fig. 16. Example of a texture flow and it's perceptual completion quality. The wavy surface appears to form a single coherent unit behind the occluders, even though there are a different number of line segments in each visible region (this effect is valid for all but the most extreme of scales). This demonstrates that orientation is distributed densely over a region, and that computing texture flows is not simply the simultaneous completion of a number of "parallel" curves.