Connection Geometry, Color, and Stereo

Ohad ben Shahar[1], Gang Li[2], and Steven W. Zucker[3]

[1] Computer Science, Ben Gurion University, Israel
ben-shahar@cs.bgu.ac.il,
[2] Real-Time Vision and Modeling Dept., Siemens Corp. Research, Princeton, NJ
gang-li@siemens.com,
[3] Computer Science, Yale University, New Haven, CT
steven.zucker@yale.edu

Abstract. The visual systems in primates are organized around orientation with a rich set of long-range horizontal connections. We abstract this from a differential-geometric perspective, and introduce the covariant derivative of frame fields as a general framework for early vision. This paper overviews our research showing how curve detection, texture, shading, color (hue), and stereo can be unified within this framework.

1 Introduction

Early vision is normally thought of as a collection of tasks, including edge detection, texture analysis, and stereo. The tools that are brought to bear to solve these tasks differ as well, with edge detection normally conceptualized in a signal-detection context, texture in the context of statistics for local image patches, and stereo as a problem in projective geometry. The integration of these tasks is normally accomplished with higher-level models. However, since the individual tasks are formulated in terms that differ from one another, these higher-level models are difficult to formulate in formal terms.

We have been pursuing a more unified approach. The motivation originally derived from our study of the visual systems in primates, especially the early cortical visual area V1. This is where orientation, as in *edge orientation*, is first abstracted, and it plays a key role in specifying the functional architecture, or layout, of cortex. While other feature dimensions, such as direction of motion or spatial frequency, are also important, for space limitations we shall not consider them here. The remaining organizing element–eye of origin–is of course necessary for stereo.

In this short paper we simply overview our work, with a focus on the geometry that runs through the early vision problems of edge detection, oriented texture and shading analysis, stereo and color. Our goal is to highlight the (differential) geometric framework that is common to all of these problems.

In the next section we introduce several concepts from modern differential geometry, especially the covariant derivative, the Frenet equations, and frame fields. Curvature emerges as the central connection between tangent orientations at nearby positions. We then illustrate the quantization of curvature for curves, which gives rise to co-circularity. Extending this to 3-D allows us to formulate the stereo problem for space curves. This provides a new framework for stereo, integrating position and orientation disparity, and illustrates how the frame-field structure can elaborate our problem formulations. The full 2-D frame field is the natural setting for oriented textures, and this involves two curvatures, one in the tangential and one in the normal directions. Finally, the extension to color, or at least hue, is sketched.

We stress that this presentation is intended as an overview of our work and as an entry to the literature. While all of the results are described in greater detail, with complete references to related research, in the pointers to the literature given as references at the end, we hope that placing these pieces in juxtaposition to one another will illustrate our confidence in — and excitement about — the frame-field representation for early vision. We recommend consulting these more complete papers for the full story.

2 Connection Geometry in the Plane

A unit length tangent vector $E(q)$ attached to point $q = (x, y) \in \Re^2$ is the natural representation of orientation in the plane. Attaching such a vector to points of interest (e.g., along a smooth curve or oriented texture) yields a unit length vector field. Assuming smoothness, an infinitesimal translation along the vector V from q yields a small rotation in the vector $E(q)$. A frame $\{E_T, E_N\}$, placed at the point q with E_T identified with $E(q)$ allows us to apply techniques from differential geometry (Fig. 1).

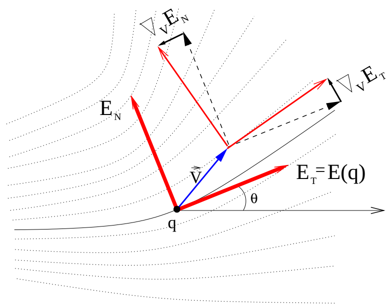

Fig. 1. Illustration of the geometry behind the connection equation and the covariant derivative.

Nearby tangents are displaced both in position and orientation according to the covariant derivatives, $\nabla_V E_T$ and $\nabla_V E_N$, which can also be represented as vectors in the basis $\{E_T, E_N\}$:

$$\begin{pmatrix} \nabla_V E_T \\ \nabla_V E_N \end{pmatrix} = \begin{bmatrix} w_{11}(V) \; w_{12}(V) \\ w_{21}(V) \; w_{22}(V) \end{bmatrix} \begin{pmatrix} E_T \\ E_N \end{pmatrix}. \quad (1)$$

Note: the 1-forms $w_{ij}(V)$ are functions of the displacement V. Since the basis $\{E_T, E_N\}$ is orthonormal, they are skew-symmetric $w_{ij}(V) = -w_{ji}(V)$. Thus $w_{11}(V) = w_{22}(V) = 0$ and the system reduces to connection equation formulated by Cartan [1]:

$$\begin{pmatrix} \nabla_V E_T \\ \nabla_V E_N \end{pmatrix} = \begin{bmatrix} 0 & w_{12}(V) \\ -w_{12}(V) & 0 \end{bmatrix} \begin{pmatrix} E_T \\ E_N \end{pmatrix}. \tag{2}$$

$w_{12}(V)$, the connection form, is linear in V, so it can be represented in terms of the frame $\{E_T, E_N\}$:

$$w_{12}(V) = w_{12}(a\ E_T + b\ E_N) = a\ w_{12}(E_T) + b\ w_{12}(E_N) \ \ .$$

giving rise to the scalars:

$$\begin{aligned} \kappa_T &\doteq w_{12}(E_T) \\ \kappa_N &\doteq w_{12}(E_N) \end{aligned} \tag{3}$$

which we interpret as tangential (κ_T) and normal (κ_N) curvatures.

Specializing to the one-dimensional case of curves, only ∇_{E_T} is necessary:

$$\begin{pmatrix} \nabla_{E_T} E_T \\ \nabla_{E_T} E_N \end{pmatrix} = \begin{bmatrix} 0 & w_{12}(E_T) \\ -w_{12}(E_T) & 0 \end{bmatrix} \begin{pmatrix} E_T \\ E_N \end{pmatrix}. \tag{4}$$

Now T,N, and κ can replace E_T,E_N, and κ_T, respectively, and this is the classical *Frenet equation* (primes denote derivatives with respect to arclength):

$$\begin{pmatrix} T' \\ N' \end{pmatrix} = \begin{bmatrix} 0 & \kappa \\ -\kappa & 0 \end{bmatrix} \begin{pmatrix} T \\ N \end{pmatrix}. \tag{5}$$

3 Curves and Co-Circularity

The original application of these ideas was to develop compatibility coefficents for a relaxation labeling process for curve detection in images; see [2]. The basic idea is that local estimates of an image curve, or approximations to the tangent, are obtained at all positions and all orientations. (These correspond to the local measurements of orientation in visual cortex of primates.) Consistent tangent estimates support one another; while inconsistent ones detract support. (This corresponds to the computation supported by the neural substrate of long-range horizontal connections.) The goal is to find that collection of tangents that maximize support. (For a technical definition of support, see [3]).

From the geometric perspective above, consistency can be interpreted directly in terms of transport along the osculating circle (a local 2-nd order approximation to the curve at each point); see Fig. 2.

4 (Oriented) Textures and Co-helicity

We now utilize the full differential geometry in the plane. In the neighborhood of an orientation within a texture, there is a full set of possible orientations; if

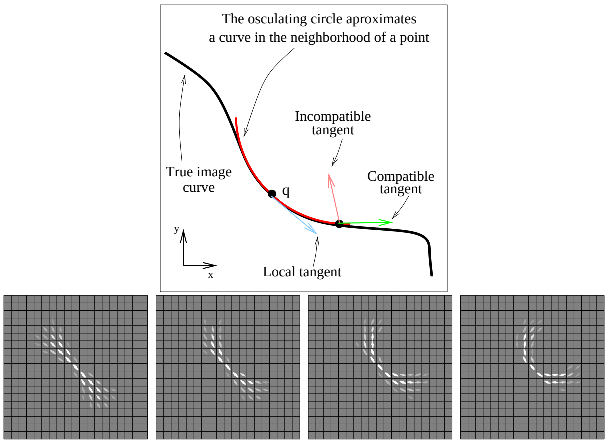

Fig. 2. The geometry of co-circularity for image curves. (top) The osculating circle provides an approximation to a curve locally via curvature. (bottom) Different quantizations of curvature indicate which tangent estimates should reinforce one another. In effect these "precompute" the different transports.

we move in any direction within the texture the orientation can change. This implies the need for compatibility functions that are fully 2-D, and these are illustrated in Fig. 3. The construction is a direct extension of co-circularity with a helicoid (in position, orientation space) as the generalization of the osculating circle for curves. Two curvatures are necessary. See [6]

5 Analysis of Hue

One normally thinks of color in (red, green, blue) coordinates. However, when the color is mapped to the psychologically more useful (intensity, hue, saturation) coordinates, the color circle emerges. Considering only hue, the value at each pixel can be represented as a vector (which points to the proper location on the hue circle). The geometry is now close to that for oriented textures, modulo π vs 2π.

Compatibility fields, which earlier were drawn among vectors, can now be drawn among hues. Notice that hue can change slightly with movement in either the tangential or the normal directions so, again, two (hue) curvatures are needed.

The resulting system can be used for denoising [4]; for segmentation; and to provide a basis for color constancy [5].

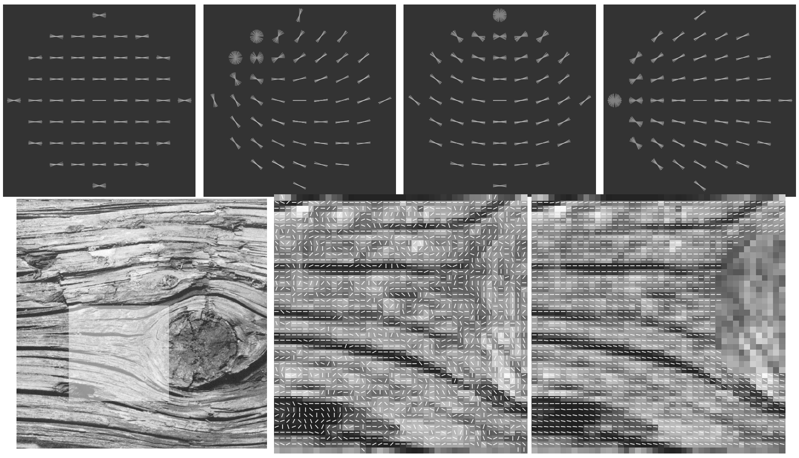

Fig. 3. Analysis of oriented textures. (top) Co-helical compatibilities for oriented textures. These are defined in a local neighborhood around the central tangent. For textures this neighborhood is 2-D. Note that now there is variation in orientation (i.e., curvature) in both the tangential and the normal directions, and that singularities arise naturally. (bottom) A Brodatz texture; initial measurements of orientation within the center-of-interest; relaxed (consistent) tangent vector field.

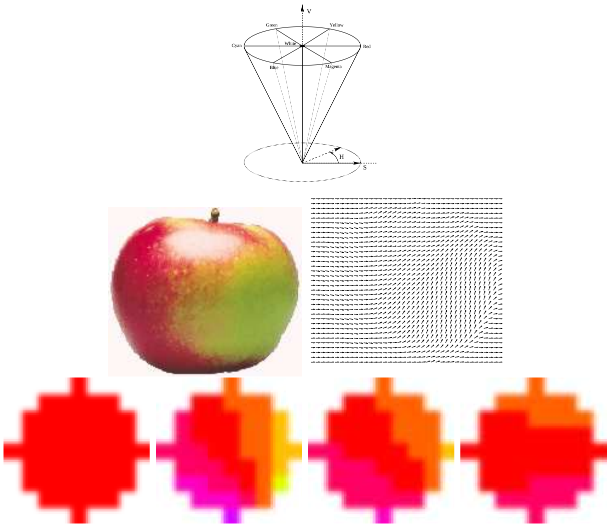

Fig. 4. Geometry of color and hue flows. (top) The intensity-hue-saturation representation for color makes the hue circle explicit. The hue at each pixel can therefore be represented as a vector, and similar geometry to that for oriented textures emerges. (middle) An apple image with the hue represented at each pixel. (bottom) Hue compatibilities now indicate how hue varies in the tangential and the normal directions around the central pixel. They can be used for noise cleaning and object segmentation.

6 Stereo

Stereo correspondence involves both differential and projective geometry. Classical projective geometry is well know in computer vision; here we stress how continuity of smooth objects (in this case curves, but also surfaces) can supplement the epipolar constraint for matching, and can supercede the ordering and other heuristic constraints.

We formulate the basic transport operation in $\Re^3$, for which we need to extend the Frenet equations to add *torsion*, or deviation from the osculating plane.

$$\begin{bmatrix} \mathbf{T}' \\ \mathbf{N}' \\ \mathbf{B}' \end{bmatrix} = \begin{bmatrix} 0 & \kappa & 0 \\ -\kappa & 0 & \tau \\ 0 & -\tau & 0 \end{bmatrix} \begin{bmatrix} \mathbf{T} \\ \mathbf{N} \\ \mathbf{B} \end{bmatrix} \tag{6}$$

The central observation is that when the standard frontal-parallel plane assumption is violated, higher-order disparities are introduced. It is these new disparities that are most useful.

To illustrate, consider the tangent component of the Frenet 3-frame. Note that this projects to a pair of (2-D) tangents, one in the left image and one in the right. They will have a classical spatial disparity as well as the higher-order *orientation disparity*; see Fig. 5. The compatibility functions and transport are defined in 3-D; and thus are implemented as relationships between *pairs* of tangents.

The stereo system for space curves is described in [7]; the extension to surfaces is in [8].

7 Summary and Conclusions

In this lecture we attempted to illustrate how the geometry of interactions for smooth objects provides a unifying theme for many of the problems in early vision. The application to the neurobiology of long-range horizontal connections in the first visual cortical area can be found in [9].

References

1. O'Neill, B., Elementary Differential Geometry, Academic Press, San Diego, 1966.
2. Parent, P., and Zucker, S.W., Trace inference, curvature consistency, and curve detection, *IEEE Trans. Pattern Analysis and Machine Intelligence*, 1989, **11**, 823 - 839.
3. Hummel, R.A., and Zucker, S.W., On the foundations of relaxation labeling processes, *IEEE Trans. Pattern Analysis and Machine Intelligence*, 1983, **PAMI-5**, 267-287.
4. Ben-Shahar, O., and Zucker, S.W., Hue fields and color curvatures: A perceptual organization approach to color image denoising, *Proc. IEEE Conf. on Computer Vision and Pattern Recognition, CVPR 03*, Madison, WI, June 2003.

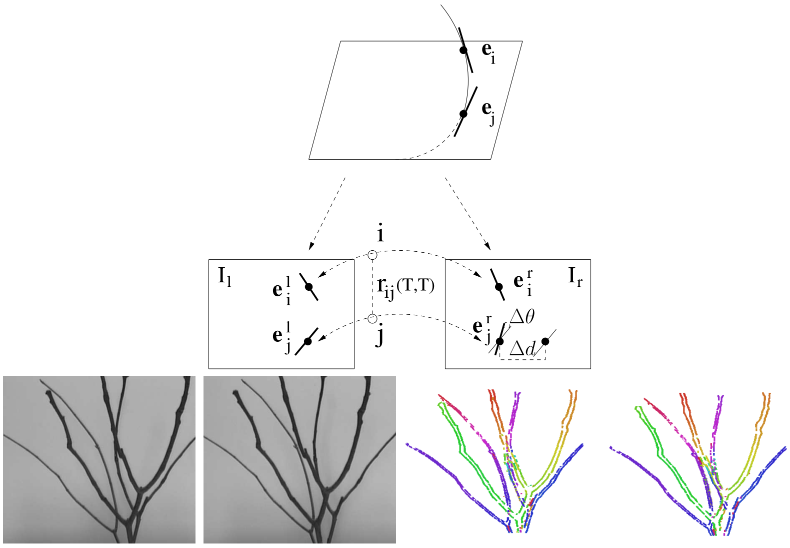

Fig. 5. The geometry of stereo. (top) A space curve in 3-D projects into the left and the right images. Shown are two (space) tangents, each of which projects to a pair of (image) tangents. Compatibilities are thus defined over pairs of tangents, and include orientation as well as positional disparities. (bottom) A complex arrangement of twigs. The left and right images are shown, as is the stereo reconstruction from two viewpoints.

5. Ben-Shahar, and Zucker, S.W., Hue geometry and horizontal connections, *Neural Networks*, 2004, **17**(5-6), 753-771.
6. Ben-Shahar, and Zucker, S.W., The Perceptual Organization of Texture Flow: A Contextual Inference Approach, *IEEE Trans. Pattern Analysis and Machine Intelligence*, 2003, **25**(4), 401 - 417.
7. Li, G., and Zucker, S.W., Contextual Inference in Contour-Based Stereo Correspondence *Int. J. of Computer Vision*, 2006, **69**(1), 59 - 75.
8. Li, G., and Zucker, S.W., Differential Geometric Consistency Extends Stereo to Curved Surfaces, *Proc. 9-th European Conference on Computer Vision*, Graz, Austria, May 7 - 13, 2006
9. Ben-Shahar, and Zucker, S.W., Geometrical computations explain projection patterns of long-range horizontal connections in visual cortex, *Neural Computation*, 2003, **16**(3), 445 - 476.