

Differential Geometry from the Frenet Point of View: Boundary Detection, Stereo, Texture and Color

Steven W. Zucker¹

ABSTRACT Frenet frames are a central construction in modern differential geometry, in which structure is described with respect to an object of interest rather than with respect to external coordinate systems. The Cartan moving frame model specifies how these frames adapt when they are transported along the object. We consider this as a model for integrating local information with information in a neighborhood for curve detection, stereo, texture, and color. These different objects results in a series of geometric compatibility constructions useful within a number of different optimization and probabilistic inference techniques.

1 Introduction

Many problems in computational vision that involve inferences over noisy, local measurements have been formulated with a geometrical component. Our goal in this Chapter is to organize a number of such problems according to their geometric content, to isolate a common thread between them that leads to differential geometry; and to introduce ideas from differential geometry to show how they can structure new approaches to seemingly unrelated computational vision problems. As described, the techniques can be used with a variety of different inference techniques, including relaxation labeling [12], belief propagation, graph cuts [5], Markov random fields, quadratic programming, and so on.

To prefigure the type of geometry we shall be concerned with, consider the problem of boundary detection. Starting with local “edge” operators that signal intensity differences in a small neighborhood around a point, the question is whether this intensity event is part of a boundary, or not. Since many objects have smooth boundaries, and since these boundaries project into the image as smooth curves, determining whether a putative

¹To appear in **Mathematical Models of Computer Vision: The Handbook** Nikos Paragios, Yunmei Chen and Olivier Faugeras (eds.), Springer, 2005

boundary point continues through an image neighborhood containing that point is often key. Mathematically, since only a neighborhood is involved, the analysis is local. Computationally, since such questions can be asked around each point in the image, the local analysis must be applicable in a neighborhood around each point; i.e., it is parallel. Differential geometry is a mathematical abstraction of boundary completion that satisfies these requirements. It will lead, as we show, to connections between the local estimates that are specialized for each problem.

Expanding the above points, recall that the best linear approximation in an infinitesimal neighborhood to a smooth (boundary) curve is its tangent, and that this tangent approximation can be made around each point. Therefore the question becomes whether nearby tangents are consistently part of a single curve. To develop an intuition about what consistent might mean, recall the classical Gestalt demonstration of perceptual *good continuation* (Fig. 1). Observe how the “Figure 8” appears to continue across the crossing point; that is, how orientation is continued along the tangent direction. Many such demonstrations were developed in the early 20th century ([16]).

Approximately a half century earlier a fundamental series of discoveries began concerning the differential geometry of curves, and they continued through the time period dominated by the Gestalt psychology movement. Frenet (in 1847) and, independently Serret (in 1851), introduced the idea of adapting a coordinate frame directly to a curve, rather than using extrinsic coordinates. The remarkable discovery was that changes in (derivatives of) this frame could be expressed directly in terms of the frame itself. The result is a beautiful expression of the theory of curves that fits precisely the requirements for perceptual organization above. The Frenet-Serret theory was extended by Darboux to surfaces a few decades later, and was then elaborated to the powerful *repère mobile*—the moving frame—by Élie Cartan. Moving frames are not slaves to any coordinate system; rather, they are adapted to the object under study, be it a curve, a surface (notice the texture flow in Fig. 1), a metric space or manifold. For computer vision applications, we shall adapt them to curves (in 2-D and in 3-D), to texture, and to color. Local approximations of how these frames move will provide the geometry of connections that can be used with the different inference techniques listed above.

There are many excellent texts describing this approach to differential geometry. We recommend [19, 24], which we have followed closely in preparing this Chapter. For related discussions see also [15]. This research was done in collaboration with Ohad Ben-Shahar, Lee Iverson, and Gang Li. I thank Pavel Dimitrov for illustrations and AFOSR, DARPA, NIH, and ONR for support.




FIGURE 1. Perceptual organization is related to Gestalt notions of “good continuation.” Observe how the “Figure 8” appears as a single curve, with smooth connections across the crossing point, and not as the non-generic arrangement of the two shapes in the middle. Such notions of orientation good continuation hold for textures as well; notice how this example appears to continue behind the occluders.

2 Introduction to Frenet-Serret

From a Newtonian perspective a curve can be thought of as the positions $\alpha(t) = (\alpha_1(t), \alpha_2(t), \alpha_3(t))$ in Euclidean 3-space swept out by a moving point α at parameter (time) t . Provided the coordinate functions $(\alpha_1, \alpha_2, \alpha_3)$ are differentiable, a *curve* can be defined as a differentiable map $\alpha : I \rightarrow \mathbb{E}^3$, from the open interval I into \mathbb{E}^3 . For now we shall assume the curve is simple, i.e., it does not cross itself, so the map is one-to-one and is an *immersion* of I into \mathbb{E}^3 .

The derivative of α gives the velocity or *tangent vector* of α at t

$$\alpha'(t) = \left(\frac{d\alpha_1}{dt}(t), \frac{d\alpha_2}{dt}(t), \frac{d\alpha_3}{dt}(t), \right)_{\alpha(t)}$$

A curve is *regular* provided these derivatives are not zero simultaneously.

A reparameterization $s = s(t)$ yields the arc-length (unit speed) parameterization in which the length of each tangent vector is 1. We denote this unit speed curve by $\beta : I \rightarrow \mathbb{E}^3$ with $\|\beta'(s)\| = 1, s \in I$.

For simplicity, we work with β for the remainder of this Section. We are interested in direction and, for non-straight lines, the rate at which the curve is bending. Intuition is helped by picturing the unit tangents as vectors in \mathbb{E}^3 attached to the points $\beta(s) \in \mathbb{E}^3$, that is, as a vector field along the curve. Euclidean coordinates for this vector field can again be differentiated:

$$\alpha''(t) = \left(\frac{d^2\alpha_1}{dt^2}(t), \frac{d^2\alpha_2}{dt^2}(t), \frac{d^2\alpha_3}{dt^2}(t), \right)_{\alpha(t)}$$

to yield the acceleration, but geometrically the following construction will be more useful. (i) Denoting the unit tangent $T = \beta'$, we obtain $T' = \beta''$, the curvature vector field. Observe T' is orthogonal to T by differentiating $T \cdot T = 1$. The direction of the curvature vector is normal to β , and its length $\kappa(s) = \|T'(s)\|, s \in I$ is the *curvature*. (ii) The vector field $N = T'/\kappa$




FIGURE 2. The Frenet frame attached to a point on a curve $\alpha(s)$ approximated to third order.

defines the *principal normal*, and (iii) the vector field $B = T \times N$ is the *binormal* vector field of β .

The *Frenet frame field* on β is the triple (T, N, B) such that $T \cdot T = N \cdot N = B \cdot B = 1$, all other dot products = 0, and the (i)–(iii) above hold (Fig. 2).

The remarkable property of this construction is that the derivatives of the frame can be expressed in terms of the frame itself. For $\kappa > 0$ and introducing the *torsion* τ we have:

$$\begin{pmatrix} T' \\ N' \\ B' \end{pmatrix} = \begin{bmatrix} 0 & \kappa & 0 \\ -\kappa & 0 & \tau \\ 0 & -\tau & 0 \end{bmatrix} \begin{pmatrix} T \\ N \\ B \end{pmatrix}. \quad (1.1)$$

These are the famous Frenet-Serret formulas. The torsion τ measures how rapidly the curve is twisting out of the (osculating) plane spanned by (T, N) . It is in this sense that the Frenet frame is adapted to the individual curve in a way that captures its essential (differential) geometric structure.

Basically all of information about the curve is contained in the Frenet-Serret formulas. The following theorem is fundamental in differential geometry: Let $\kappa, \tau : I \rightarrow \mathbb{R}$ be continuous ($\kappa(s) > 0, s \in I$). Then there is a curve $\beta : I \rightarrow \mathbb{E}^3$ with curvature function $\kappa(s)$ and torsion $\tau(s)$. Any two such curves differ only by a proper Euclidean motion.

Writing the Taylor approximation to the curve in the neighborhood of $\beta(0)$, and then substituting the Frenet formulas above and keeping only the dominant terms, we obtain:

$$\beta(s) \approx \beta(0) + s\beta'(0) + \frac{s^2}{2}\beta''(0) + \frac{s^3}{6}s\beta'''(0) \quad (1.2)$$

$$\approx \beta(0) + sT_0 + \kappa_0 \frac{s^2}{2}N_0 + \kappa_0\tau_0 \frac{s^3}{6}B_0. \quad (1.3)$$




FIGURE 3. Two ways to think about the local structure of a curve in the plane. (left) The Frenet Frame is a (tangent, normal) coordinate frame that is adapted to the local structure of each point along a curve; and (right) the osculating circle is that circle with the largest contact with the curve among all circles tangent at that point.

Thus the Frenet approximation shows how the tangent, curvature, and torsion effect the curve at each point (Fig.2).

3 Co-Circularity in $\mathbb{R}^2 \times S^1$

We now focus on curves in the plane \mathbb{E}^2 . Observe that the first two terms in the Frenet approximation give the line in which the tangent (or best linear approximation) lies; the first three terms give the best quadratic approximation (a parabola) which, expressed in the (x,y) plane, has the shape $y = \kappa_0 x^2/2$ near $\beta(0)$.

The quadratic approximation around a point is determined by the curvature at that point, which can be defined in another way. Suppose the curve is not straight, and choose any three points on β in the neighborhood of $\beta(0)$. Taking the limit as the three points approach $\beta(0)$, the *osculating circle* at that point is obtained. This is the unique circle tangent to the curve at that point such that its center lies on the normal and its radius is the inverse of the curvature (Fig.3).

The quadratic parabola is approximated by the *osculating circle* at that point, an observation introduced for the geometry of co-circularity [20]². The basic idea is illustrated in Fig. 4, which shows how local measurements of the tangent to a curve at an arbitrary point q and at a nearby point in its neighborhood have different orientations. The geometry of consistency is given by Frenet: if the frame in the neighborhood of q is transported

²Because of space limitations, references are very limited; we recommend that the original publications are consulted for additional references.

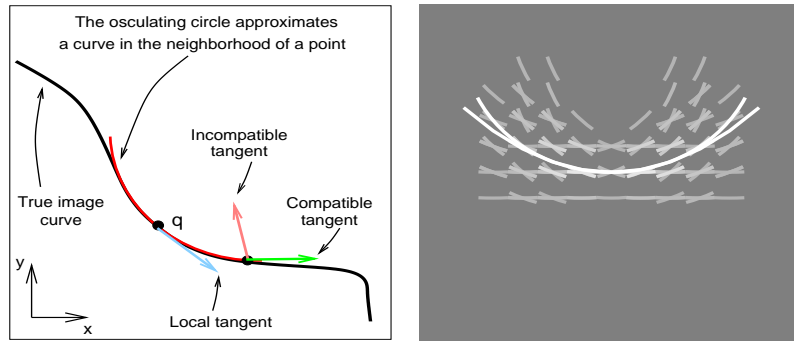


FIGURE 4. The geometry of co-circularity for curve detection in images. (left) Measurements of orientation differ at points along a curve. To determine whether they are consistent, nearby tangents are transported along the osculating circle approximation to the curve. If the transported tangents agree they are consistent; otherwise not. (right) To accomplish this transport operation in images, tangent position, orientation, and curvature must be discretized. This shows those nearby tangents that are consistent with a horizontal tangent at the center; that is, those tangent which, if transported along a (discretized) approximation to the osculating circle would support the central, horizontal tangent. (The width of the curve for this example is taken to be 3 pixels.) In the language of relaxation labeling, this is called an excitatory compatibility field. Note that the osculating circle and parabola approximations agree to within a fraction of a pixel over this neighborhood.

along the curve to q , it should match the frame at q . If it does not, it is inconsistent.

However, the curve must be known before transport can be applied, but this is what we seek. The solution to this chicken-and-egg problem is to transport not along the actual curve, but along its approximation. We earlier showed that curvature dictates this approximation, and it can either be measured directly (which is what we think happens in neurobiology, [9]) or estimated by other means ([2]). In any case, once the system is discretized, the osculating circle and parabolic approximations agree to within a fraction of a pixel over the neighborhoods involved (Fig. 4); cf. [13]. Such geometric compatibility fields can be used with a number of different inference techniques, including relaxation labeling [12], belief propagation, and Bayes [14]. They are related to the forms that arise in elastica [18, 10]. For a different attempt to minimize a functional in curvature, see [23].

3.1 Multiple Orientations and Product Spaces

Thus far in this Chapter we have been concentrating exclusively on simple, regular curves. But the “figure 8” example in Fig. 1 is not simple, and it provided the motivation for the geometric approach. Which way should

the curve be continued at the crossing point? For such examples, although $\beta(s_1) = \beta(s_2)$ for $s_1 \neq s_2$ at the crossing point, we have $\beta'(s_1) \neq \beta'(s_2)$, which provides a clue. Instead of assuming there is only one unique tangent per pixel, which is commonplace in computer vision [7], we shall allow more than one.

To allow multiple tangents at each position, it is natural to attach a copy of the space of all possible tangents to each position (Fig. 3.1). Since in principle tangent angle is distributed around the circle and position is a real number, the resultant space is $\mathbb{R}^2 \times S^1$. (Note differences from the classical coordinate representation.) This space is an example of another fundamental construct in modern differential geometry, the *unit tangent bundle* associated with a surface in \mathbb{E}^3 . Intuitively one might think of a surface as being covered by (i.e., as a union of) all possible curves on that surface. More generally, the tangent bundle to a surface is the union of tangent spaces at all points. If the surface is 2-D, the tangent bundle is 4-D. The geometric compatibility fields can be applied in parallel to all tangents in this space. (We will be generalizing this construct in the next few Sections, and will show examples then.)

[26] discusses the relevance of this product construction for the neurobiology of vision.




FIGURE 5. The need for higher-dimensional spaces than the image arises in representing non-simple or piecewise-regular curves. Since *a priori* a curve could be passing through any pixel at any orientation, it is natural to represent the (discretized) circle (the space of all unit vectors) S^1 at each (discretized) position (left). When the non-simple “figure 8” is lifted into the resultant space, the lift is a simple curve in $\mathbb{R}^2 \times S^1$ (right). The (position, orientation) space, which is abstract from the image, is sufficient to represent all possible curves in the image.




FIGURE 6. (a) Cartoon of the stereo relaxation process. A pair of space tangents associated with the Frenet approximation around the point with tangent \mathbf{e}_j . Each of these tangents projects to a (left,right) image tangent pair; compatibility between the space tangents thus corresponds to compatibility over (left,right) image tangent pairs. The projected tangents are shown as thick lines. One left image tangent is redrawn in the right image (as a thin line) to illustrate positional disparity (Δd) and orientation disparity ($\Delta\theta$). The compatibility between the tangent pair (i) and the pair (j) is denoted r_{ij} . Of course, for the full system the complete Frenet 2-frames are used to infer the Frenet 3-frame attached to the space curve. (b) Just as the osculating circle provided a local model for transport for image curves, a section of a helix provides a local model for a space curve. The (T, N) components of the Frenet 3-frame define the osculating plane, which rotates as the frame is moved along the space curve.

4 Stereo: Inferring Frenet 3-Frames from 2-Frames

We now move to 3-space, and consider the problem of inferring the structure of space curves from projection into two images. Earlier we showed that a curve in \mathbb{R}^3 has a tangent, normal, and binormal Frenet frame associated with every regular point along it. To sketch a geometric approach to stereo compatibility, for simplicity consider only the tangent in this frame and imagine it as an (infinitely) short line segment. This space tangent projects into a planar tangent in the left image and a planar tangent in the right image. Thus, space tangents project to pairs of image tangents. Now, consider the next point along the space curve; it too has a tangent, which projects to another pair of image tangents, one in the left image and one in the right image. Thus, in general, transport of the Frenet 3-frame

in \mathbb{R}^3 from the second point back to the first has a correspondence in the left-right image pairs of 2-frames. [17] have developed this transport idea to find corresponding pairs of image tangents such that their image properties match, as closely as the geometry can be approximated, the actual space tangents (Fig. 6). They show, in particular, that the stereo projection operator can be inverted to give the Frenet 3-frame and the curvature, but not the torsion. This builds on the related work of [8, 22, 21]

Two notions of disparity arise from the above transport model. First, the standard notion of positional disparity corresponds, through the camera model, to depth. Second, an orientation disparity is introduced if the space tangent is not in the epipolar plane. In the computational vision literature, orientation disparity is largely unexplored, but it is widely studied in visual psychophysics [11]. The geometric viewpoint shows how to use position and orientation disparities together. Typical reconstructions from this algorithm are shown in Fig. 7.




FIGURE 7. 3D reconstruction of a twig pair.(a) Left image (b) Right image; note in the highlighted region that subtleties in using the ordering constraint arise. Furthermore, occlusion of branches gives rise to discontinuities in orientation. Representing such discontinuities as multiple tangents facilitates proper matching. (c) Reconstruction. Depth scale is shown at right (units: meters).

5 Covariant Derivatives, Oriented Textures, and Color

We now denote orientation in the plane as a unit length tangent vector $\hat{\mathbf{E}}(\vec{\mathbf{q}})$ attached to point $\vec{\mathbf{q}} = (x, y) \in \mathbb{R}^2$. With such tangent vectors attached to every point of an oriented texture results in a unit length vector field, which creates a need to generalize the notion of transport: the frame can now be moved in any direction in the texture, rather than just along the curve.

Assuming good continuation as in the Introduction, a small translation $\vec{\mathbf{V}}$ from the point $\vec{\mathbf{q}}$ should rotate the vector $\hat{\mathbf{E}}(\vec{\mathbf{q}})$ a small amount. Following the Frenet model, the frame $\{\hat{\mathbf{E}}_T, \hat{\mathbf{E}}_N\}$ is placed at the point $\vec{\mathbf{q}}$ and the basis vector $\hat{\mathbf{E}}_T$ is identified with $\hat{\mathbf{E}}(\vec{\mathbf{q}})$ – the tangent vector at $\vec{\mathbf{q}}$ (Fig. 5). Note

that $\hat{\mathbf{E}}_T$ is drawn at an angle θ – the local orientation measured relative to the x-axis – such that $(\hat{\mathbf{q}}, \theta) \in \mathbb{R}^2 \times S^1$. Nearby tangents are displaced both in position and orientation according to the *covariant derivatives*, a tensor object whose components are essentially the partial derivatives of the underlying pattern. (Covariant derivatives generalize the earlier derivatives which were taken only along the curve; i.e., with respect to the arc length parameter s .) For vector fields the covariant derivative is taken in a direction given by another vector field, and is a vector. Again following Frenet, we observe that such covariant derivatives, $\nabla_{\vec{\mathbf{V}}}\hat{\mathbf{E}}_T$ and $\nabla_{\vec{\mathbf{V}}}\hat{\mathbf{E}}_N$, are naturally represented as vectors in the basis $\{\hat{\mathbf{E}}_T, \hat{\mathbf{E}}_N\}$ itself:

$$\begin{pmatrix} \nabla_{\vec{\mathbf{V}}}\hat{\mathbf{E}}_T \\ \nabla_{\vec{\mathbf{V}}}\hat{\mathbf{E}}_N \end{pmatrix} = \begin{bmatrix} w_{11}(\vec{\mathbf{V}}) & w_{12}(\vec{\mathbf{V}}) \\ w_{21}(\vec{\mathbf{V}}) & w_{22}(\vec{\mathbf{V}}) \end{bmatrix} \begin{pmatrix} \hat{\mathbf{E}}_T \\ \hat{\mathbf{E}}_N \end{pmatrix}. \quad (1.4)$$

The coefficients $w_{ij}(\vec{\mathbf{V}})$ are *1-forms*, real-valued functions defined on tangent vectors. They are functions of the displacement direction vector $\vec{\mathbf{V}}$, and since the basis $\{\hat{\mathbf{E}}_T, \hat{\mathbf{E}}_N\}$ is orthonormal, they are skew-symmetric $w_{ij}(\vec{\mathbf{V}}) = -w_{ji}(\vec{\mathbf{V}})$. Thus $w_{11}(\vec{\mathbf{V}}) = w_{22}(\vec{\mathbf{V}}) = 0$ and the system reduces to:

$$\begin{pmatrix} \nabla_{\vec{\mathbf{V}}}\hat{\mathbf{E}}_T \\ \nabla_{\vec{\mathbf{V}}}\hat{\mathbf{E}}_N \end{pmatrix} = \begin{bmatrix} 0 & w_{12}(\vec{\mathbf{V}}) \\ -w_{12}(\vec{\mathbf{V}}) & 0 \end{bmatrix} \begin{pmatrix} \hat{\mathbf{E}}_T \\ \hat{\mathbf{E}}_N \end{pmatrix}. \quad (1.5)$$

This begins to resemble the Frenet-Serret formulas but is more general; it is *Cartan's connection equation*; $w_{12}(\vec{\mathbf{V}})$ is called the *connection form*. Since $w_{12}(\vec{\mathbf{V}})$ is linear in $\vec{\mathbf{V}}$, it can be represented in terms of $\{\hat{\mathbf{E}}_T, \hat{\mathbf{E}}_N\}$:

$$w_{12}(\vec{\mathbf{V}}) = w_{12}(a \hat{\mathbf{E}}_T + b \hat{\mathbf{E}}_N) = a w_{12}(\hat{\mathbf{E}}_T) + b w_{12}(\hat{\mathbf{E}}_N) .$$

The relationship between nearby tangents is thus governed by two scalars at each point.

$$\begin{aligned} \kappa_T &\triangleq w_{12}(\hat{\mathbf{E}}_T) \\ \kappa_N &\triangleq w_{12}(\hat{\mathbf{E}}_N) \end{aligned} \quad (1.6)$$

We interpret them as *tangential* (κ_T) and *normal* (κ_N) curvatures, since they represent a directional rate of change of orientation in the tangential and normal directions, respectively.

The connection equation describes the local behavior of orientation for the general two dimensional case, but it can be specialized to the one-dimensional case of curves developed earlier by considering only $\nabla_{\hat{\mathbf{E}}_T}$:

$$\begin{pmatrix} \nabla_{\hat{\mathbf{E}}_T}\hat{\mathbf{E}}_T \\ \nabla_{\hat{\mathbf{E}}_T}\hat{\mathbf{E}}_N \end{pmatrix} = \begin{bmatrix} 0 & w_{12}(\hat{\mathbf{E}}_T) \\ -w_{12}(\hat{\mathbf{E}}_T) & 0 \end{bmatrix} \begin{pmatrix} \hat{\mathbf{E}}_T \\ \hat{\mathbf{E}}_N \end{pmatrix}. \quad (1.7)$$

which, in our earlier notion, becomes:

$$\begin{pmatrix} T' \\ N' \end{pmatrix} = \begin{bmatrix} 0 & \kappa \\ -\kappa & 0 \end{bmatrix} \begin{pmatrix} T \\ N \end{pmatrix}. \quad (1.8)$$

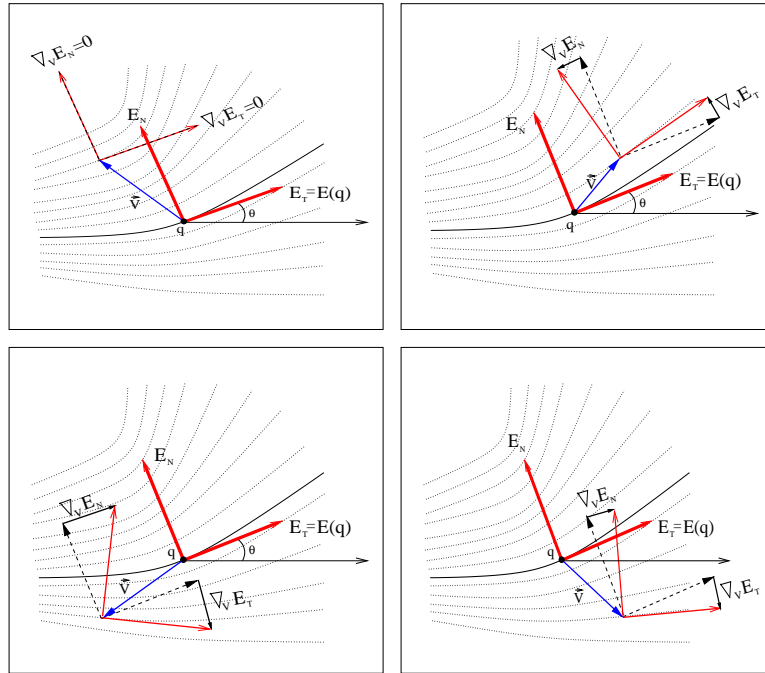


FIGURE 8. Displacement (transport) of a Frenet frame within a vector field or an oriented texture amounts to rotation, but differs for different displacements. The covariant derivative specifies the frame's initial rate of rotation for any direction vector \vec{V} . The four different cases in this figure illustrate how this rotation depends on \vec{V} both quantitatively (i.e., different magnitudes of rotation) and qualitatively (i.e., clockwise, counter-clockwise, or zero rotation). A pure displacement in the tangential direction (\hat{E}_T) specifies one rotation component (the tangential curvature) and a pure displacement in the normal direction (\hat{E}_N) specifies the other (normal curvature) component.

We refer to κ_T as the *tangential curvature* and κ_N as the *normal curvature* - they represent the rate of change of the dominant orientation of the texture flow in the tangential and normal directions, respectively. In the language of frame fields, κ_T and κ_N are just the *coordinate functions* of $\nabla\theta$ with respect to $\{E_T, E_N\}$.

In the case of curves, the theory of frames is coupled to ordinary differential equations. For vector fields and texture flows, partial differential equations arise. In particular, since E_T and E_N are rigidly coupled, and we have

$$\begin{aligned}\kappa_T &= \nabla \times E_T \\ \kappa_N &= \nabla \cdot E_T \quad .\end{aligned}\tag{1.9}$$

If κ_T and κ_N were known functions of position $q = (x, y)$, a PDE could be solved for the rotation angle $\theta(q)$. Thus κ_T and κ_N are not completely

independent, and integrability conditions arise. In particular, unless κ_T and κ_N are both equal to zero, they cannot be constant simultaneously in a neighborhood around q , however small, or else the induced flow is nonintegrable. [3] show that, given any texture flow $\{E_T, E_N\}$, its curvature functions κ_T and κ_N must satisfy the relationship

$$\nabla \kappa_T \cdot E_N - \nabla \kappa_N \cdot E_T = \kappa_T^2 + \kappa_N^2$$




FIGURE 9. Texture compatibility fields are discretizations of a helicoid approximation to a flow lifted into $\mathbb{R}^2 \times S^1$. Three examples are shown: **(A)** both curvatures are zero; this is the analog to a straight line for curves; **(B)** tangential curvature is zero and normal curvature is positive; this shows a local portion of a texture flow in which the integral curves converge to a (singular) point, as lines converge to a point in the distance; and **(C)** both the tangential and the normal curvatures are positive. This is the general case: notice how singular points (where all orientations are possible) arise. These are indicated as multiple line segments displayed at the same position.

With osculating circles the natural local model for the geometry of regular planar curves, and helices the natural model for regular space curves, [3] show that the natural local model for textures and flows is a helicoid in $\mathbb{R}^2 \times S^1$. This follows intuitively because each streamline or integral curve through the flow can be locally approximated by a section of an osculating circle; this lifts to a section of a helix. The helicoid is a ruled surface built of these lifts. Local sections of the helicoid can be projected into the image and discretized to provide connection or compatibility fields for textures and flows (Fig. 9).

The result of applying this system to overlapping flows is shown in Fig. 10. Notice in particular how woven textures can be thought of as multiple threads, or curves, overlapping one another. This emerged from our discussion of representing multiple orientations at each point. When overlapping textures are lifted into $\mathbb{R}^2 \times S^1$ their structure separates just as the “figure 8” separated at the crossing point. But now, in a discrete sense, such multiple values are very common.

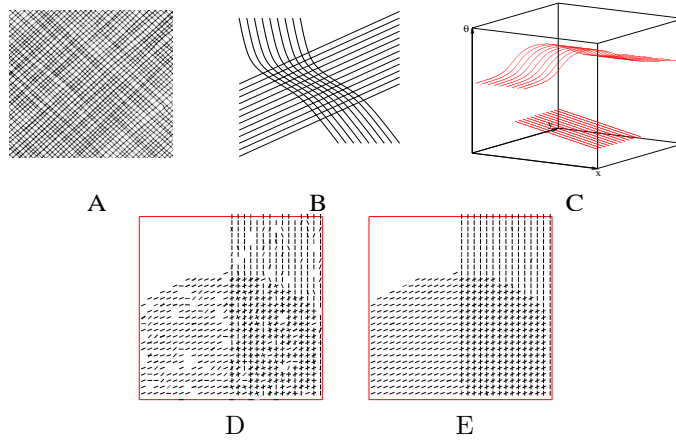


FIGURE 10. Examples of texture patterns rich in orientation. (A) A woven texture with two dominant orientations. This is an extension of (B) two overlapping textures, which are naturally separated when lifted into $\mathbb{R}^2 \times S^1$ in (C). The bottom panels illustrate how a noisy pattern (D) is refined using the geometric compatibilities in Fig. 9 to (E), thereby enforcing a Gestalt-like good continuation of the flows.

5.1 Hue Flows

While color is normally thought of as a point in (R,G,B)-space, it can also be represented in the psychophysically motivated HSV color space. Here a color image is a mapping $\mathcal{C} : \mathbb{R}^2 \rightarrow \mathcal{S}^1 \times [0, 1]^2$ (see Fig.11). The hue component across the image is a mapping $\mathcal{H} : \mathbb{R}^2 \rightarrow \mathcal{S}^1$ and thus can be represented as a unit length vector field over the image, which [3] called the *hue field*. Displays of the hue field reveal that it may vary greatly, albeit smoothly, even *within* perceptually coherent objects (see Fig 12).

Many color image enhancement algorithms are based on a form of anisotropic diffusion [1, 6], using either a vectorial representation or a manifold representation [25]. While diffusion in color space can work within very smooth regions, it does have the tendency to blur inappropriately.

Hue compatibility fields can be defined analogously to texture compatibility fields—see[4]. As expected, concepts of hue curvatures naturally arise, which express how the hue is flowing from one image position to those in its neighborhood. Just as with texture flows, a tangential and a normal hue curvature are required. Since the local behavior of the hue is characterized (up to Euclidean transformation) by this pair of curvatures, it is natural to conclude that nearby measurements of hue should relate to each other based on these curvatures. Or, put differently, measuring a particular curvature pair at a point should induce a field of coherent measurements, i.e., a hue function in its neighborhood. Coherence of hue to its spatial context




FIGURE 11. The HSV color representation in $\mathcal{S}^1 \times [0, 1]^2$ and the color wheel.

can then be determined by examining how well it fits with those around it. Again, a helicoidal approximation in (position, hue) space arises.

Such flows are relevant to image denoising; for estimating mutual reflectance and color bleeding; for estimating smooth surface variations as separate from lighting variations (for lightness algorithms); and for separating cast shadow boundaries and highlights from other types of intensity edges.

6 Discussion

In this Chapter we co-developed ideas from modern differential geometry and problems in computer vision. The differential geometry was based on Frenet and Serret's ideas of attaching frames directly to curves, rather than expressing curve structure in terms of extrinsic coordinate functions. Such ideas were carried to a remarkable stage by Cartan, whose moving frame concept is now central in mathematics. The covariant derivative emerges for differential variation of frames in flows, as the normal derivative was useful for transporting a frame along a curve.

The moving frame concept provides a natural abstraction for perceptual organization problems, at least for those that can be defined over short distances. We considered curve detection in 2D and stereo as the projection of 3D curves to illustrate the power of this geometric abstraction. Techniques for integrating orientation disparity with positional disparity emerged. But the real power was seen for flows, in which textures and hues were considered.

Although the notion of tangent was introduced as the best linear approximation to a curve, modern definitions abstract via a limiting operation to an equivalence class of curves. Our discussion attempted to avoid any unnecessary abstraction, so that all concepts had a direct counterpart in computer vision terms.




FIGURE 12. A flow perspective on color images is provided by their hue fields. These are typically piecewise smooth. Most importantly, hue can vary smoothly even within perceptually coherent objects. **(top)** A natural image of an apple with varying hue. Notice that the everyday expression of “red apple” is limited. The corresponding hue field changes smoothly across the image of the apple’s surface. **(bottom)** A 3D representation of the hue field, where hue is represented as height. Identifying the top face with the bottom (since hue is a circle) leads to the (position, hue) space.

Consideration of non-simple curves motivated an elaboration of the types of representations normally considered in computer vision from image-based ones to those that attach a space of possibilities at each point. It is commonplace to assume boundaries have a well-defined orientation at each point, but this holds for only a restricted class of curves. Local occlusion clues involving “T” junctions provide an important example of non-smooth curves, and our elaborated representation is capable of handling them as well.

The space of possible frames also has an important representation in differential geometry, and is related to fibre bundles. We just touched on such concepts in this Chapter, but fully expect them to be playing a much richer role in future applications of differential geometry to computational vision.

7 REFERENCES

- [1] S. Acton. Multigrid anisotropic diffusion. *IEEE Transactions on Image Processing*, 7(3):280–291, 1998.

- [2] J. August and S. Zucker. Sketches with curvature: The curve indicator random field and markov processes. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 25:387–401, 2003.
- [3] O. Ben-Shahar and S. Zucker. The perceptual organization of texture flow: A contextual inference approach. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 25:401–417, 2003.
- [4] O. Ben-Shahar and S. Zucker. Hue geometry and horizontal connections. *Neural Networks*, 17:753–771, 2004.
- [5] Y. Boykov, O. Veksler, and R. Zabih. Fast Approximate Energy Minimization via Graph Cuts. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 23:1222–1239, 2001.
- [6] V. Caselles, B. Coll, and J.-M. Morel. Geometry and color in natural images. *Journal Mathematical Imaging and Vision*, 16:89–105, 2002.
- [7] R. Deriche. Using canny’s criteria to derive a recursively implemented optimal edge detector. *International Journal of Computer Vision*, 1(2), MAY 1987.
- [8] F. Devernay and O. D. Faugeras. Computing differential properties of 3-d shapes from stereoscopic images without 3-d models. In *Proc. IEEE Conf. on Computer Vision and Pattern Recognition*, 1994.
- [9] A. Dobbins, S. W. Zucker, and M. S. Cynader. Endstopping and curvature. *Vision Research*, 29:1371–1387, 1989.
- [10] B. Horn. The curve of least energy. *ACM Trans. Mathematical Software*, 9:442–460, 1983.
- [11] I. P. Howard and B. J. Rogers. *Binocular Vision and Stereopsis*. Oxford Univ. Press, 1995.
- [12] R. Hummel and S. W. Zucker. On the foundations of relaxation labeling processes. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 6:267–287, 1983.
- [13] B. Kimia, I. Frankel, and A. Popescu. Euler spiral for shape completion. In K. Boyer and S. Sarker, editors, *Perceptual Organization for Artificial Vision Systems*, pages 289–310. Kluwer Academic Publishers, 2000.
- [14] D. C. Knill and W. Richards. *Perception as Bayesian Inference*. Cambridge University Press, 1996.
- [15] J. J. Koenderink. *Solid Shape*. MIT Press, Cambridge, Massachusetts, 1990.

- [16] K. Koffka. *Principles of Gestalt Psychology*. Harcourt, Brace, and World, New York, 1935.
- [17] G. Li and S. Zucker. A differential geometric approach to stereo correspondence. In *Second IEEE Workshop on Variational, Geometric, and Level Set Methods in Computer Vision*, 2003.
- [18] D. Mumford. Elastica and computer vision. In C. Bajaj, editor, *Algebraic Geometry and Its Applications*, pages 507–518. Springer Verlag, 1993.
- [19] B. O’Neill. *Elementary Differential Geometry*. Academic Press, London, 1966.
- [20] P. Parent and S. W. Zucker. Trace inference, curvature consistency and curve detection. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 11(8):823–839, August 1989.
- [21] S. B. Pollard, J. E. W. Mayhew, and J. P. Frisby. Pmf: A stereo correspondence algorithm using a disparity gradient limit. *Perception*, 14:449–470, 1985.
- [22] C. Schmid and A. Zisserman. The geometry and matching of lines and curves over multiple views. *International Journal of Computer Vision*, 40(3):199–233, 2000.
- [23] A. Sha’ashua and S. Ullman. Structural saliency: The detection of globally salient structures using a locally connected network. In *International Conference on Computer Vision*, pages 321–327, 1988.
- [24] M. Spivak. *A Comprehensive Introduction to Differential Geometry*. Publish or Perish, Houston, 1975.
- [25] B. Tang, G. Sapiro, and V. Caselles. Color image enhancement via chromaticity diffusion. *IEEE Transactions on Image Processing*, 10:701–707, 2001.
- [26] S. W. Zucker. Which computation runs in visual cortical columns? In J. L. van Hemmen and T. Sejnowski, editors, *Problems in Systems Neuroscience*. Oxford University Press, 2004.