Complexity and Confusion in Computational Vision

Steven W. Zucker
Depts. of Computer Science and Electrical Engineering
Yale University
New Haven, CT 06520-8285

Abstract

We propose a geometrical perspective toward early vision. Edge elements, we maintain, should be viewed as a local approximation to curves, or as tangents, and raises the question of which curves can arise. The Whitney theorem classifies these, and suggests separate analyses for the tangential and the normal photometrical variations. Grouping the tangent elements together then raises a local-to-global transition. Problems in this transition can be viewed as follows for an edge element inferred from an unknown image. Is this local edge part of a long curve, or perhaps part of a texture? If the former, which is the next element along the curve. If the latter, is the texture like a hair pattern, in which nearby elements are oriented similarly, or like a spaghetti pattern? Are there other generic possibilities? The questions raise issues of complexity and of dimensionality, since curves are 1-D and textures are 2-D. It also raises questions of confusion: when is a curve-like representation appropriate? We propose a measure of representational complexity that seeks to answer these questions. A consequence of our analysis is the requirement for intermediate representations based on tangent maps, which supports the initial approach to edge detection, and leads to a new approach to early segmentation.

1 Introduction

The definition of an edge, one of the foundations on which much of visual processing is built, remains unsettled. The usual one, which has roots back to Ernst Mach in the mid-19^{th} century, is that edges are bright-to-dark transitions (or vice-versa), and that these transitions can be detected from measurements provided by (discrete approximations to) linear differential operators. Mach, for example, preferred the Laplacian operator, and Marr has revived modern interest in it. However, this view has turned out to be problematic, and much frustration in visual shape analysis derives from the unsatisfactory nature of early edge processing. In brief, the edges returned by standard tools do not correspond in any systematic way to boundaries of objects, and it is these boundaries–not bright-to-dark transitions–that are needed for visual shape analysis (see Fig. 1). Our goal in this paper is to propose a geometric foundation for early vision in an attempt to get right to the matter of what comprises a boundary (Fig. 2). Our proposal is based on concepts from differential geometry and geometric measure theory. In the process of achieving this goal, we shall overview our recent research in edge detection, and shall concentrate on the geometric problems associated with interpreting the role an edge might be playing in an image. The geometric approach in the end leads us to conclude that one should think not in terms of edges, which are image-defined constructs, but rather in terms of tangents, which are differential-geometric constructs. It is this abstraction that reveals the connection back to boundaries.

2 Boundaries and their Detection

What, precisely, is a boundary? Mach's intuition was founded in photometry, and is based on the observation that, since the reflectance of an object is typically different from the background, the object will project to an image region with different intensity from those regions to which the background projects. This intuition breaks down, however, when the object has internal structure.

To articulate this internal boundary structure, consider a smooth surface, such as a ball, that doesn't have any. The surface of the ball has a 2-dimensional tangent plane everywhere, and this tangent plane projects smoothly into the image except along the locus of positions at which the viewing vector just grazes the ball; at these positions the tangent plane becomes singular, and collapses to a single (one-dimensional) tangent. That it to say, the tangent plane folds away from the line of sight. An integral curve through these tangents defines the boundary of the ball. Thus, insofar as continuous mathematical approximations hold, each tangent defines a proper edge element.

To see that edge elements can arise from other configurations, consider a bulge of cloth. Edges can now arise

(a)

(b)

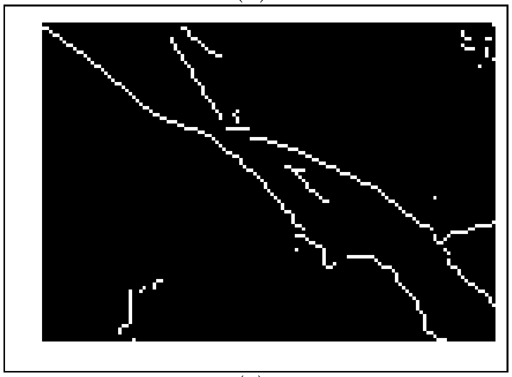

(c)

Figure 1: An illustration of the geometric problems in interpreting standard edge maps. In this case the edge map (b) is obtained from the Canny operator, Matlab implementation, scale=3. While this looks sufficent at first glance, a detail from the shoulder/arm region in (c) shows the incorrect topology; neither the detail of the shoulder musculature, nor the separation from the arm, are localized correctly. A second look at the full edge map in (b) confirms the incorrect topology. Examine, for example, the left shoulder, neck, chin configuration. More generally, consider which of these transitions are exterior boundaries, or interior ones. Under which circumstances should boundaries end?

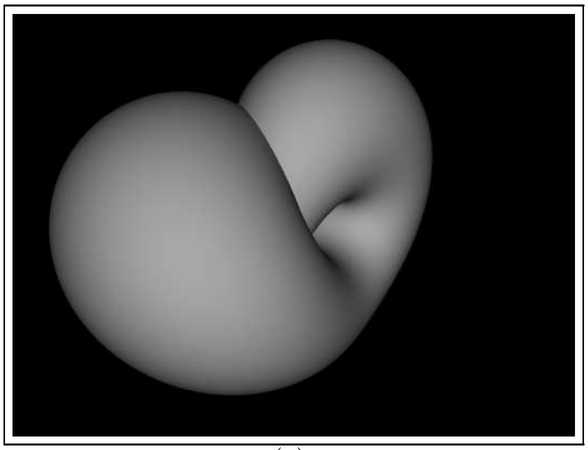

(a)

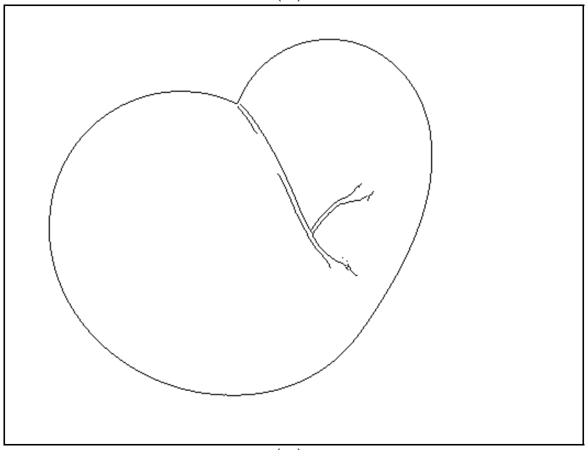

(b)

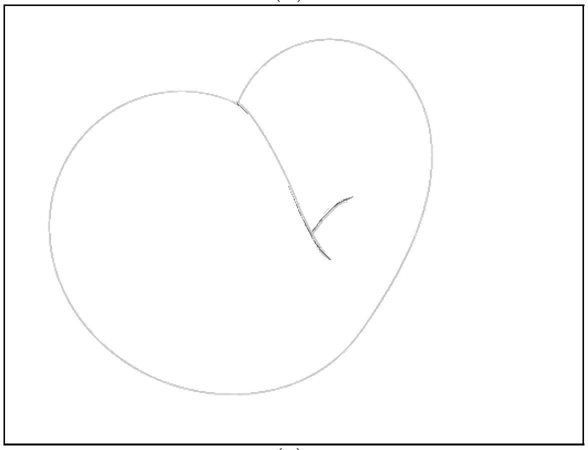

(c)

Figure 2: The motivation from differential topology for early vision. The image of the Klein bottle in (a) shows how "T" junctions can arise from occlusion relationships (e.g., at the top of the figure), and how certain interior edges can end (e.g., where the fold smoothly joins the body). The Canny edge structure shown in (b) is inconsistent with both of these topological observations. Notice how the boundary "T"-junction is not connected, how it smooths the outline, and how the interior folds blur into the shading. In (c) is shown the output of the logical/linear operator. Notice how the "T"-junctions are maintained, and how the contours end at cusps. Dark tangents signify lines; gray ones edges; see text.

from within the cloth, with the tangent plane disappearing along the top of the bulge and then re-appearing on the other side. The tangent plane transition across a bulge (from cloth-to-cloth) is thus different from that across the ball, because it never re-appears. Moreover, this type of fold can end within the body of the cloth, as in the wrinkes around one's shirt sleeve or the muscles in one's shoulder; see Fig. 2 and Fig. 3b.

2.1 Whitney's Classification Theorem

Whitney has classified the generic maps from smooth surfaces onto smooth surfaces, and has shown that these two situations are the only ones that can occur generically. Generically means that the configuration does not change with small changes in viewpoint. He referred to them as the fold and the cusp (the position where the cloth fold disappears). Furthermore, by considering a view of a mountain range it is easy to observe that generic boundaries are not globally smooth; rather, they are punctuated by discontinuities. Thus we have three observations regarding boundaries:

1. *local edge elements arise in different contexts* such as exterior boundaries or interior folds. Textures, of course, provide further examples of complex contexts for edge elements.
2. *discontinuities should be expected in bounding contours* from occlusion relationships. They will often have the form of T-junctions.
3. *interior fold elements can end* within a smooth surface.

2.2 Normal Intensity Configurations and Logical/Linear Operators

We begin with the first observation to provide formal and computational support for the idea that local edge elements can be interpreted as tangents to contours. The first advantage is that, with the tangent direction established, one can examine the intensity profiles in the tangent and the normal directions separately. Note that most other edge operators average these together, although not necessarily uniformly. Thus we observe immediately that,

- *Normal direction:* The fold condition in Whitney's Theorem takes on a different intensity profile for a bounding edge (which involves a dark-to-light transition) than that from an interior fold (which involves a light-to-dark-to-light transition) or vice versa. This latter profile is often called a line, and must be separated from the former. Standard linear operators blur both together. Thus is a common photometric confusion in edge detection.
- *Tangential direction:* The differential interpretation demands that continuity conditions exist (that is, that the limit of one point approaching another must exist). This corresponds to continuity contrainst on the intensity pattern.

A necessary condition for a tangent to exist is that the above intensity and continuity conditions must be satisfied. We have developed a class of non-linear local operators, called logical/linear operators, that use Boolean conditions to test whether the above structural criteria are met. "Edge operators" are separated from "line operators", and lines can arise either in light-dark-light conditions (typical of a crack or a crease) or dark-light-dark conditions (typical of a highlight). Note that both of these latter conditions refer to surface markings, rather than to surface boundaries. The edge map in Fig. 1(c) contains both bright-to-dark transitions, shown in gray, and bright-dark-bright transitions shown in black.

The tangent interpretation also leads to a solution for representing orientation discontinuities, but requires a more modern view of a discontinuity. While it is classically the case that no unique tangent exists at a non-regular point on a curve, it is also the case that multiple tangents can be defined to exist there. Informally, taking a limit into the point of discontinuity from one side yields one tangent; a limit from the other side yields another. The two tangents span a 2-dimensional space, the Zariski tangent space to the curve; it is precisely this difference in dimensionality between the tangent spaces at regular and singular points that we exploit. The logical/linear operators are arranged into columns, so that multiple values of tangent orientation are possible at each point. Those points at which multiple tangents are established are the discontinuities.

Code for a system that implements the three logical/linear operators [1] in columns, one for edges, one for bright lines, and one for dark lines, is available from via anonymous ftp from cvc.yale.edu. Researchers are invited to experiment with this system. Results are shown in Fig. 3.

2.3 Edge Maps, Tangent Maps, and Grouping

The result of our logical/linear edge operator [1] at a given scale is shown in Fig. 3b on the image of the statue ("Paolina", see Fig. 3a), and raises the following observations. For the shoulder region (Fig. 3c,d), the underlying object is simple and a curve representation seems appropriate to group the edge elements. If we examine instead regions subtending part of the hair structure (Fig. 3e,f), then choosing a curve representation and walking along a hair would lead very quickly to confusion, since it will be difficult to know on which part of the curve one is. A texture representation in

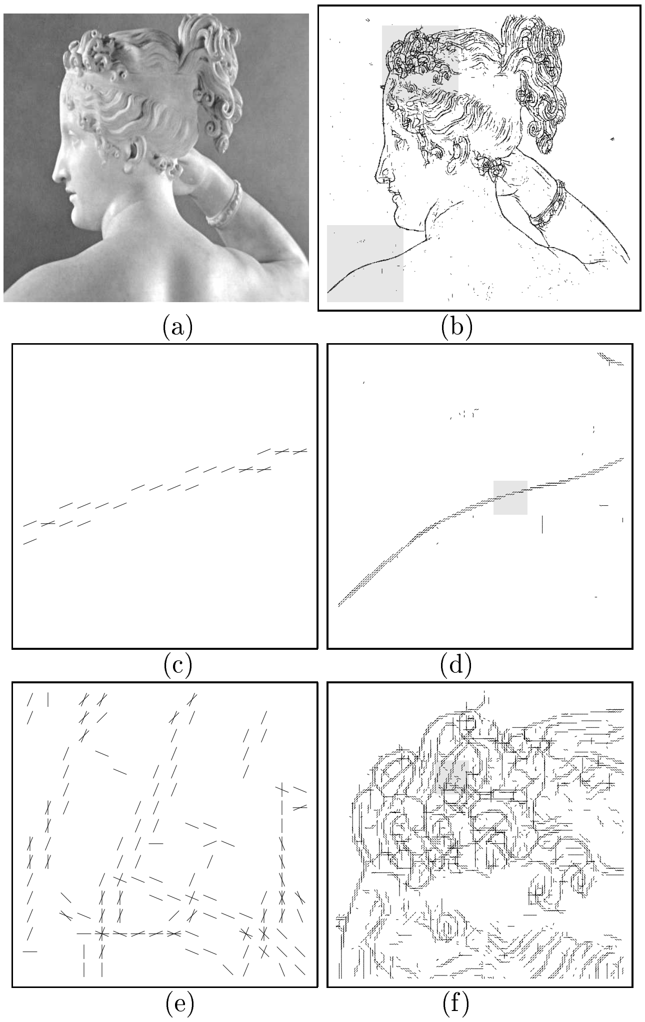

Figure 3: The output of logical/linear edge operators on the Paolina image. Notice the proper fold structure around the upper right shoulder, the arm, and the neck. T-junctions are clear, and topology maintained. This figure also illustrates the subtlety of "walking through" a tangent map. (b) shows the two regions highlighted in gray. The left shoulder region, viewed in increasing detail (d,c), intuitively appears curve-like, while the hair region, (f,e) is not. How should this hair region be characterized? Furthermore, are multiple tangent values at a position caused by quantization effects, or by crossing hairs? Is it possible to separate these cases. Notice how this difference between quantization and transversality is important for interpreting early descriptions.

this case seems more appropriate. The remainder of the paper will deal with these issues.

Although we have been speaking of the the logical/linear operators as if they signal tangents, actually they only return a distribution of positions at which the above intensity signatures (namely, bright-dark, and bright-dark-bright transitions) are obtained. This is a very different notion than that of tangent. As we show in the next Section, however, there is a geometric-measure-theoretic definition of tangent that provides some degree of unification.

3 Continuous rectifiable curves and curve-like sets

This section is an attempt to answer the question: "what is a curve?" in the context of curve detection. The most common definition of a curve is the one of Jordan, namely that a curve Γ is the range of a continuous map α from an interval I to Euclidean space (typically $\mathbf{R}^2$ or $\mathbf{R}^3$). Two other basic notions usually follow in any study of elementary differential geometry: the one of *length*, denoted $L(\Gamma)$, and the one of best linear local approximation, namely the *tangent* to Γ at $x = \alpha(t)$, denoted $T(x)$. Within these are embedded the notions of local representation (the tangent) and global measure (the length), which are tied together through the map α. Unfortunately, however, the map α is not given for general vision problems, and more abstract notions of a curve are thus required. This is why a generalization of the previous ideas through measure theory is necessary.

3.1 Geometric Measure Theory

One way to compute the length, area or volume of an object is to use the Hausdorff s-dimensional measure $\mathcal{H}^s$, where in the case of a smooth rectifiable curve, $s = 1$, in the case of a surface, $s = 2$. Consider the problem of defining the length of a set E in the plane (henceforth we shall concentrate on the case $s = 1$). Hausdorff's construction was to cover the set with small circles, and to take the sum of the diameters. If the balls are restricted to be smaller than some given value $\delta > 0$, and if the 'most economical' covering is chosen, we get an approximation of the length of the set at resolution δ. Allowing arbitrary covers, instead of covers by balls, gives us an outer measure, and for $\delta > 0$ we write

$$\mathcal{H}_\delta(E) = \inf \sum_i |U_i|$$

where $|U|$ is the *diameter* of U, (i.e., $|U| = \sup\{|x - y| : x, y \in U\}$) and $\{U_i\}$ is any sequence of sets of diameter less than δ covering E. The infimum here is taken over all (countable) δ-cover $\{U_i\}$ of E and we can easily show that $\mathcal{H}_\delta(E)$ increases as δ decreases, therefore:

Definition 1 (Hausdorff measure [2]) *The* one dimensional Hausdorff measure *of E is given by*

$$\mathcal{H}(E) = \lim_{\delta \to 0} \mathcal{H}_\delta(E) = \sup_{\delta > 0} \mathcal{H}_\delta(E)$$

It is an important result that, if Γ is a curve, then the Hausdorff measure $\mathcal{H}(\Gamma)$ and the Jordan length $L(\Gamma)$ coincide [2]. But most importantly, we have

Definition 2 (Curve-like set [2]) *A set E with $0 < \mathcal{H}(E) < \infty$, will be called a* curve-like set.

3.1.1 Local structure of curve-like sets

Before discussing the existence of tangents for curve-like sets, we will present an alternate definition of a tangent that doesn't rely on a parameterization of the set. This definition is due to Besicovitch:

Definition 3 (Tangent: Besicovitch [3]) *A curve-like set E has a* tangent $T_B(x)$ *at x in the direction $\pm\theta$ if*

1. *x "is on" the set E,*
2. *for every angle $\phi > 0$,*

$$\lim_{r \to 0} \frac{\mathcal{H}(E \cap (B_r(x) \backslash S_r(x, \theta, \phi) \backslash S_r(x, -\theta, \phi)))}{r} = 0$$

where $B_r(x)$ is the ball of radius r centered at x and $S_r(x, \theta, \phi)$ is the sector of radius r at angle θ with opening ϕ.

Suppose $x \in E$, then this definition means that at x the set E is locally concentrated around the line $T_B(x)$ with orientation θ passing through x. The first condition in this definition can be formally expressed in terms of the density of the set E at x [2]. The second condition ensures that the concentration is around the tangent line only. Fig. 4b illustrates the definition, namely that the second condition consists in looking at the rate of growth of what is found outside the local angular sector centered at x. If this rate of growth is much faster than r, the curve is ensured (from the first condition) to be concentrated around the line with orientation θ at x.

How does this definition of tangent relate to the usual definition of a tangent to a parameterized curve? In his seminal work, Besicovitch [3] showed that his new parameterization free definition was indeed equivalent to the usual definition. From now on we will therefore write $T(x)$ for the Besicovitch tangent.

The standard Besicovitch definition doesn't allow however for multiple tangents at a point, a key requirement in the type of sets we are studying when doing curve detection [4]. As we saw earlier, in real world situations objects occlude each other leading to discontinuities in bounding contours and to T-junctions. At these points "multiple tangents" must be represented. The following is an extension of the Besicovitch tangent to allow the representation of multiple tangents at a point:

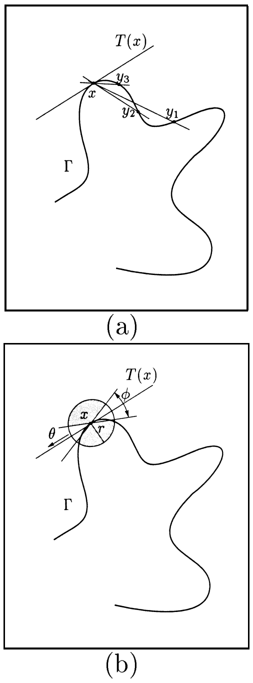

Figure 4: This figure illustrates in (a) the intuitive definition of a tangent T to a curve Γ at a point $x = \gamma(t)$. Take a sequence $\{y_1, y_2, \cdots\}$ of points on the curve converging to x. Draw the lines passing through y_i and x. The "limit line" gives us the tangent $T(x)$ to the curve at x. In (b) we illustrate the parametrization-free definition of the tangent, looking at a cone that shrinks around the point x (see text).

Definition 4 (Multiple tangents) *A curve-like set E has a* tangent set $\Theta(x)$ *at x if x "is on" the set E and for every angle $\phi > 0$,*

$$\lim_{r \to 0} \frac{\mathcal{H}(E \cap (B_r(x) \backslash (\bigcup [S_r(x, \theta, \phi) \cup S_r(x, -\theta, \phi)])))}{r} = 0$$

but also, for each $\theta \in \Theta(x)$, $\exists\, r_0$ and ϕ_0 such that $\forall 0 < \phi < \phi_0$ and $0 < r < r_0$,

$$\mathcal{H}(E \cap (B_r(x) \cap S_r(x, \theta, \phi) \cap S_r(x, -\theta, \phi))) > 0$$

The union $\bigcup$ is over $\theta \in \Theta(x)$. As in the definition of the Besicovitch tangent, the "on the set" condition can be formally expressed in terms of densities [5]. The condition given by the first equation in the definition prevents things to be too crumpled around the point, while the final inequation ensures that indeed there is something going on in the directions contained in the tangent set.

Existence of multiple tangents It is easy to build a set with multiple tangents. The graph of the absolute value function $G_f = \{(y, f(y)) : f(y) = |y|\}$ for instance has a multiple tangent at $x = (0, 0)$. In this case $\Theta(x) = \{\pi/4, 3\pi/4\}$. In both cases we have for all $r > 0$ and

$0 < \phi < \pi/8$

$$\mathcal{H}(E \cap (B_r(x) \cap S_r(x, \theta, \phi) \cap S_r(x, -\theta, \phi))) = r > 0,$$

By bundling the non-empty tangent sets, we obtain our key intermediate representation, the *tangent map*:

Definition 5 (Tangent map) *Given a curve-like set E, we define the* tangent map *τ to be*

$$\tau = \bigcup_{x \in E} (x, \Theta(x)),$$

where $\Theta(x)$ is the set of tangents (in the Besicovitch sense) at x.

When the set E is rectifiable, it can be proved that multiple tangents cannot occur 'too often'.

4 From edge detection to discrete curve-like sets

We are now ready to establish the connection back to the opening sections of this paper, where logical/linear operators where shown to provide a dense collection of positions at which tangents should be signaled. The main feature of the Besicovitch tangent was the dense collection of points within an elongated cone. This is exactly like what was done in edge detection except that now it is the density of pixels supporting a particular contrast distribution. There is however another major difference: the concept of finite scale. In the continuous domain, objects can be classified into those with finite and with infinite length. Those with finite length were called curve-like sets. In edge detection, the distinction will be more subtle. Since the operators will be band limited, the inferred objects will be finite but it is in some sense the "degree of finiteness" that matters and will lead a partitioning scheme for discrete curve-like sets.

4.1 From edge detection to discrete curve-like sets

When inferring curve-like objects from images, one is confronted with issues such as discretization, quantization and choice of scale. We will assume here that the scale σ at which the tangents are detected is fixed. The result of the inference will be:

Definition 6 (Discrete tangent map) *Given the responses of the linear/logical edge/line detectors over a grey level image I, the* discrete tangent map *$\hat{\tau}$ will be:*

$$\hat{\tau} = \{((x_i, y_i), \hat{\theta}) | \hat{\theta} \in \hat{\Theta}(x_i, y_i)\}$$

$i = 1, 2, \cdots$, where $\hat{\theta}$ is a quantized orientation at the discrete image coordinate (x_i, y_i).

Any entry in the map in fact corresponds to an equivalence class of curves passing through the cell (x_i, y_i) with orientation $\hat{\theta}$. Relaxation labeling and co-circularity constraints can be used to ensure some degree of noise removal.

This approximate position and orientation at scale σ will have to be considered when assessing the complexity and integrating data. In particular, as mentioned earlier, one big issue is the problem of transversality versus quantization. If at a particular position there exists more than one tangent, can it be decided whether these correspond to a quantization artifact or to the fact that two lines are actually crossing at that position? It is possible to answer this question using the following complexity measures.

5 Mapping complexity

Section 3 introduced continuous curve-like sets and their local structure. The last section made the link between edge detection and the theory in the continuous domain using the local structure of discrete curve-like sets, the discrete tangent map. In this section we will refine this classification and provide a means of deciding which representation should be used for the integration process.

5.1 From dilations to complexity

To assess the complexity of the tangent map, we will use a variant of an approach due to Minkowski. Originally it consisted in covering the object with balls of radius ϵ and computing the measure of the dilated object. The rate of growth of this measure can be linked to the complexity of the set. Our approach differs in the way that the dilation is achieved, in particular it will not be done isotropically. After reviewing the standard technique, we present our variation.

5.1.1 Isotropic dilations

Minkowski dilations are routinely used in mathematical morphology [6] and for the estimation of fractal dimension [2]. The approach consists in creating a new set which is the Minkowski sum with a dilating (structuring) element. The dilation is done isotropically over the set. More formally, given a subset $X = \mathbf{R}^2$, we have

Definition 7 (Minkowski sum, dilation [6]) *The* Minkowski sum *is defined in the power set of X by putting*

$$A \oplus B = \{a + b, a \in A, b \in B\}.$$

The translate *of B by z being $B_z = B \oplus \{z\}$, the set $\{z : A \cap B_z \neq \emptyset\}$ of the points z such that A hits the translate B_z is called the* dilation *of A by B and is denoted $A \oplus \breve{B}$.*

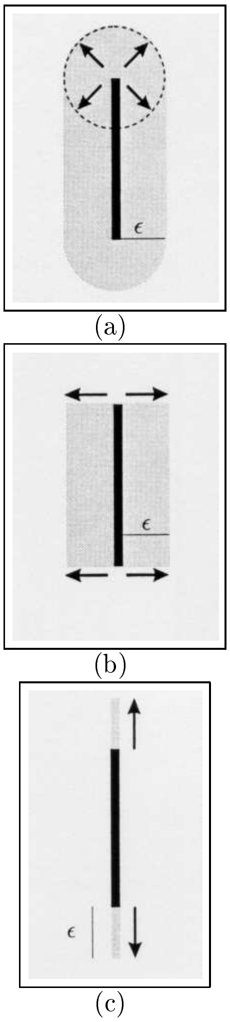

Figure 5: Isotropic and oriented dilations. (a) isotropic dilation with a ball of radius ϵ. (b) normal dilation (c) tangential dilation. Oriented dilations are possible because of the intermediate representation provided by the Besicovitch tangent sets.

An example of the isotropic dilation of a line segment with a ball is shown on Fig. 5a. In fractal analysis, the case of a dilation with a ball is often called the *Minkowski sausage.* Dilation with other shapes (squares, segments, etc.) are called *generalized Minkowski sausages.*

5.1.2 Oriented dilations

One of the key differences between our approach and the standard fractal analysis techniques is that the dilations will not be done isotropically but will adapt to the local structure of the set. We call them *oriented dilations* and we will show that they are useful for separating sets of different complexity. The local structure, of course, will depend on interpreting the logical/linear operators as signaling the tangent to a geometrical object.

Again, taking the example of the line segment, Fig. 5b,c illustrate the concepts of both normal and tangential dilations:

Definition 8 (Normal and tangential dilations) *Let E be a curve-like set and τ its tangent map. The* normal dilation $E_N(\epsilon)$ *of E at a scale ϵ is the dilation of the set E with the segment $(-\epsilon, \epsilon)$ in the direction normal to the tangents $\theta \in \Theta(x)$ at x (Fig. 5b). The* tangential dilation $E_T(\epsilon)$ *is obtained by dilating the segment $(-\epsilon, \epsilon)$ in the direction of the tangent (Fig. 5c).*

Key observation 1 *The departure from the standard Minkowski dilation approach by using oriented dilations will be essential for our analysis since it will segregate the classification* curve *vs.* texture *(using normal dilations) from the one of* dust *vs.* curve *(using tangential dilations).*

5.1.3 Normal and tangential complexity

The local information contained in the tangent map τ can be used to calculate what we will call the *normal complexity* $C_N(\epsilon)$ and the *tangential complexity* $C_T(\epsilon)$ of a curve-like set at a given scale ϵ. The main idea is to be able to look at the rate of growth of the measure of the dilated sets. In the case of normal complexity, it will be $|E_N(\epsilon)|_2$ at scale ϵ. If the area of $E_N(\epsilon)$, denoted $|\cdot|_2$, is of order α, we say that the *normal complexity* $C_N(\epsilon)$ for E is $2-\alpha$:

Definition 9 (Normal complexity) *Let E be a curve-like set. If $E_N(\epsilon)$ is the normal dilation of E at scale ϵ, and if $|\cdot|_2$ denotes its area, then the* normal complexity log-log plot *is defined as follows*

$$\left(\log\left(\frac{1}{\epsilon}\right), \log\left(\frac{1}{\epsilon^2}|E_N(\epsilon)|_2\right)\right) \tag{1}$$

moreover the normal complexity $C_N(\epsilon)$ *at scale ϵ will be the left derivative of the normal complexity log-log plot.*

From the structure of curve-like sets, one can show that the normal complexity is indeed well-defined.

Much in the same way as we did for the normal complexity, we will build a measure of complexity from tangential dilations and call it *tangential complexity.* The tangential complexity will be sensitive to ends of lines, corners, and points of high curvature; and it will allow segregation between dust and curves and scale selection for integration.

Definition 10 (Tangential complexity) *Let E be a curve-like set. If $E_T(\epsilon)$ is the tangential dilation of E at scale ϵ, and if $|E_T(\epsilon)|_{p,1}$ denotes its perimeter, then the* tangential complexity log-log plot *is defined as follows*

$$\left(\log\left(\frac{1}{\epsilon}\right), \log\left(\frac{1}{\epsilon}|E_T(\epsilon)|_{(p,1)}\right)\right) \tag{2}$$

moreover the tangential complexity $C_T(\epsilon)$ *at scale* ϵ *will be the left derivative of the tangential complexity log-log plot.*

An important remark needs to be placed at this point and for this we will use a simple example. Let E be a single line or a set of lines with the same orientation. Then dilate the set tangentially and then look at the rate of growth of measure of the dilated set. For instance, a line of length l, becomes a line of length $l + 2\epsilon$ when dilated tangentially, therefore we have

$$|E_T(\epsilon)|_{(p,1)} = 2(l + 2\epsilon)$$

Notice here that the 'measure' chosen is the *length* (not the area as was the case in the normal complexity). Therefore, we expect the rate of growth, α to be between 0 and 1. We defined $C_T(\epsilon)$, the tangential complexity, to be $1 - \alpha$. In our example, when l is large with respect to ϵ, the rate of growth is small and α is "close" to zero, therefore it has a "curve" structure. Otherwise, if l is small with respect to ϵ, then α is closer to 1 and the object would be better described as having a "dust-like" structure. Although it is easy to show that on simple examples the definition is well behaved, for general curve-like sets, the validity of the tangential complexity still remains a conjecture. In the light of our results, it will become evident however that such a measure of complexity can be most valuable to the integration process.

5.2 Complexity indexes and complexity map

The complexity measures presented in the last sections had the characteristics of being local in scale. This was a major departure from the standard 'fractal analysis' approach in which the rate of growth is studied around zero. The other key difference is to make the complexity measure local in space by computing it over a given compact region $\Omega(x)$ centered at x. The normal complexity at x can be obtained by first restricting the dilation to the region $\Omega(x)$:

$$E_N(x, \epsilon) = E_N(\epsilon) \cap \Omega(x). \tag{3}$$

The case of the tangential dilation is a little more complex but basically consists in not dilating any further than the region $\Omega(x)$, therefore leading to $E_T(x, \epsilon)$

Definition 11 (Complexity indexes) *The* normal and tangential complexity indexes *at* x *over a region* $\Omega(x)$ *are denoted* $C_N(x, \epsilon)$ *and* $C_T(x, \epsilon)$ *and are obtained by looking at the rate of growth of their local oriented dilations. More formally they are obtained as the left derivatives of their corresponding log-log plots*

$$\left(\log\left(\frac{1}{\epsilon}\right), \log\left(\frac{1}{\epsilon^2}|E_N(x,\epsilon)|_2\right)\right) \tag{4}$$

for $C_N(x, \epsilon)$*, and*

$$\left(\log\left(\frac{1}{\epsilon}\right), \log\left(\frac{1}{\epsilon}|E_T(x,\epsilon)|_{(p,1)}\right)\right) \tag{5}$$

for $C_T(x, \epsilon)$.

In a similar fashion as we did in the case of the tangent map, we can bundle all the complexity indexes and build our main tool, the *complexity map*:

Definition 12 (Complexity map) *Given a curve-like set* E*, we define the* complexity map C *to be*

$$C = \bigcup_{x \in E} (x, C_N(x, \epsilon), C_T(x, \epsilon)).$$

5.2.1 Discrete complexity map

The discrete complexity map will be obtained by computing locally normal and tangential complexity at each image position subtending a non-empty tangent. The algorithm we used consisted into projecting the unit tangents from the discrete tangent map and then dilating the resulting intermediate image. Computing the area of the dilated set and estimating the rate of growth from log-log data give us two numbers, the normal and tangential complexity indexes, to be bundled into the complexity map.

5.3 Choosing the appropriate representation

How can we build from a geometric measure point of view a classification scheme for the types of objects (texture, oriented textures, curves, etc.) that can occur in curve detection? How can the complexity map be used as an index of representation? The partition of the space of discrete curve-like sets can be done as follows, recognizing four different categories of objects:

1. a *dust-like* curve-like set: the tangent map is sparse, and the object almost nowhere extends along its length in the local neighborhood;
2. a *curve* curve-like set: those places where a curve representation is completely adequate. The object extends along its length and the normal density is low almost everywhere along it in a local neighborhood;
3. *turbulent* curve-like set: characterized by objects that do not extend along their length but are dense in their normal direction;
4. a *texture flow* curve-like set: where the object extends along its length but the density in the normal direction is high.

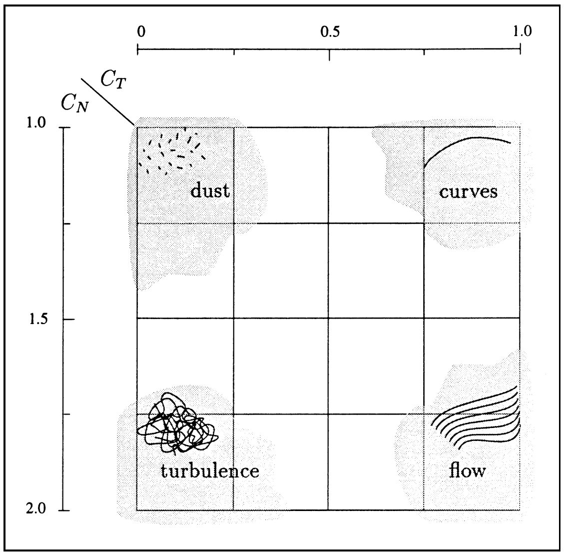

Figure 6: THE COMPLEXITY SPACE. The tiled square represents the space of complexity index pairs that can be encountered. The normal complexity varies between 1.0 and 2.0 while the tangential complexity is between 0.0 and 1.0. Each of the corners of the complexity space corresponds to a different type of curve-like substructure. Partitioning the complexity space will result in an image segmentation scheme bound to the structure of the objects in the visual scene.

The space of valid tangential/normal complexity pairs is a subset of $\mathbf{R}^2$, namely $[0.0, 1.0] \times [1.0, 2.0]$. This is because the tangential complexity is a number between 0.0 and 1.0 and the normal complexity a number between 1.0 and 2.0. Segmenting this space allows us to partition the space of possible patterns into equivalence classes. The following segmentation gives us the following nomenclature:

- low normal complexity - low tangential complexity : *dust*
- low normal complexity - high tangential complexity: *curve*
- high normal complexity - low tangential complexity: *turbulence*
- high normal complexity - high tangential complexity: *flow*

See Figure 6.

The four patterns listed in the introduction arise as the extrema in normal and tangential complexity space.

5.4 Applying oriented dilations

One way understand the implications of the oriented dilations is to look at the results when applied to the Paolina tangent map. In this particular example, the dilation will be done at two scales ϵ_1 and ϵ_2. The result will display the dilated sets with two different grays: the lighter one for the smallest scale and the union of the darker and the lighter set corresponds to the dilation at the largest scale. Moreover, in all cases the projection of the tangent map has been overlaid on the resulting dilated sets. In Fig. 7 we showed the result of tangential and normal dilations on the hair and shoulder regions for the tangent map of the Paolina image. The chosen scales were $\epsilon_1 = 0.1$ and $\epsilon_2 = 0.05$ for this example. From these images we get a better feel at the underlying idea for the complexity. For the shoulder region (Fig. 7a,b), the normal dilated set seems to grow linearly while the tangential dilations don't grow at all after some point. On the hair area (Fig. 7c,d), it's the normally dilated regions that stop growing at an early stage, while it is more difficult to describe at a quick glance what is happening to the rate of growth of the tangential dilated sets.

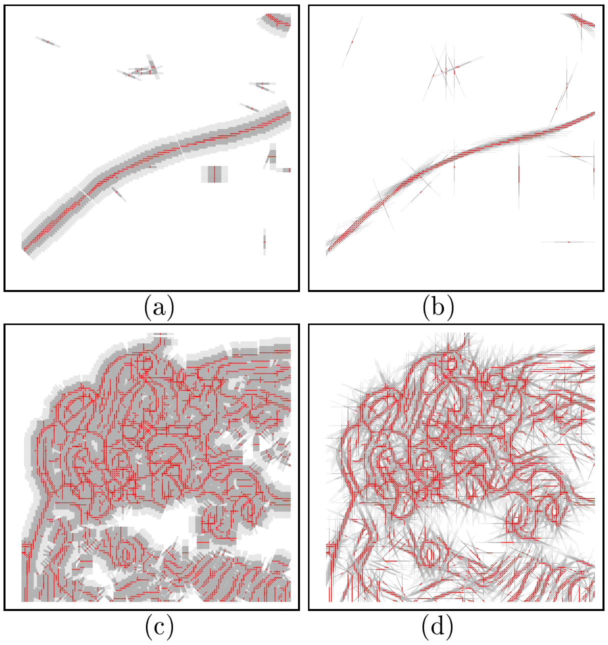

Figure 7: Close-ups on the oriented dilated sets for the Paolina image. In (a) and (c), the normal dilation, while in (b) and (d), the tangential dilation.

If we apply these techniques to the Paolina tangent map, the resulting segmentation is shown in Fig. 8. Notice how some of the bounding contours of the head are given by the curve-like set, while others are given by the turbulence set. It is classically assumed in computer vision that boundaries are always nice curves; to obtain such curves for hair-like boundaries, significant image-based smoothing is invoked. This introduces an artifact of scale into boundary processing, and blurs the differences between "hair" and "chin". We maintain that, as supported by this more complex representation, the natural complexity of hair regions should be bounded by turbulence.

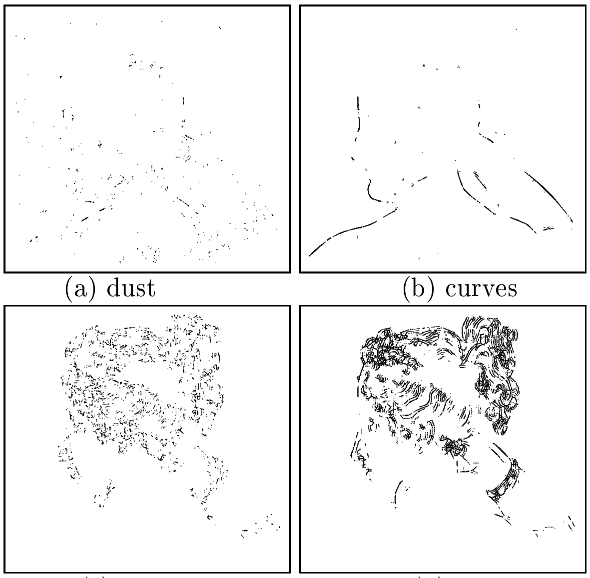

Figure 8: Segmented image of the Paolina image using the complexity map and our classification rule. The integration under a 'curve' representation should only take place within the "curves" substructure.

6 Conclusions

We have overviewed two different aspects of early vision, the first stage of edge detection and a second stage of segmentation. Processing in both stages was based on an abstraction to the level of differential topology and geometry, with the particular abstraction that edge elements should be thought of as tangents. This abstraction was supported by the Whitney theorem about maps between smooth surfaces; between transversal occlusions between surfaces, and by the Besicovitch construction from geometric measure theory. It led to complexity measures applicable to the map of tangents.

The key advantage at both levels was the local coordinate frame provided by the tangents. Edge detection was facilitated by differential processing of intensities in the tangential and the normal directions, which was related to the contrast signature along and across Whitney folds. Tangent map segmentation was facilitated by differential processing of tangents, which led to a separation between dust, curves, texture flows, and turbulence.

More generally, this intermediate level of abstraction provides the foundation on which shape analysis can be performed.

Acknowledgments

Research supported by AFOSR, NSF, NSERC, and Yale University. Much of this material was developed in collaboration with Benoit Dubuc (the complexity theory) and Lee Iverson (the logical/linear operators). Patrick Huggins computed the Klein bottle figures.

References

[1] Lee A. Iverson and Steven W. Zucker. Logical/linear operators for image curves. *IEEE PAMI*, 1995.

[2] J. K. Falconer. *Fractal Geometry: Mathematical Foundations and Applications.* Wiley, 1990.

[3] A. S. Besicovitch. On the fundamental geometrical properties of linearly measurable plane sets of points. *Mathematische Annalen*, 98:422–464, 1928.

[4] S. W. Zucker, A. Dobbins, and L. Iverson. Two stages of curve detection suggest two styles of visual computation. *Neural Computation*, 1:68–81, 1989.

[5] Benoit Dubuc and Steven W. Zucker. Indexing visual representations through the complexity map. In: *Fifth International Conference on Computer Vision*, 1995.

[6] G. Matheron. *Random Sets and Integral Geometry.* Wiley, 1975.