Contour Fragment Grouping and Shared, Simple Occluders

Jonas August
Yale University
Dept. of Electrical Engineering

Kaleem Siddiqi
McGill University
Dept. of Computer Science

Steven W. Zucker
Yale University
Dept. of Computer Science

Submitted to CVIU

Please send correspondence to:
Jonas August
Yale University
P.O. Box 208285, Yale Station
New Haven, CT 06520-8285
Phone: (203) 432-1290
FAX: (203) 432-0593
jonas.august@yale.edu

Abstract

Bounding contours of physical objects are often fragmented by other occluding objects. Long-distance perceptual grouping seeks to join fragments belonging to the same object. Approaches to grouping based on invariants assume objects are in restricted classes, while those based on minimal energy continuations assume a shape for the missing contours and require this shape to drive the grouping process. While these assumptions may be appropriate for certain specific tasks or when contour gaps are small, in general occlusion can give rise to large gaps, and thus long-distance contour fragment grouping is a different type of perceptual organization problem. We propose the long-distance principle that those fragments should be grouped whose fragmentation could have arisen from a shared, simple occluder. The gap skeleton is introduced as a representation of this virtual occluder, and an algorithm for computing it is given. Finally, we show that a view of the virtual occluder as a disc can be interpreted as an equivalence class of curves interpolating the fragment endpoints.

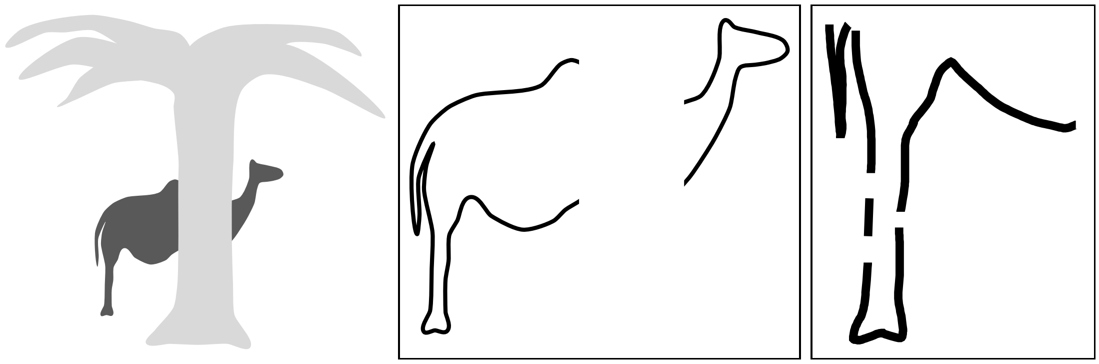

Figure 1: Different distance scales for contour fragmentation. (LEFT) The bounding contour of a camel is broken by a foreground palm tree. (CENTER) Curve fragments remaining after depth separation using T-junctions. This is long-scale fragmentation. (RIGHT) Magnification of rear leg. Observe slight contour gaps can be caused by sensor noise. This is short-scale fragmentation. The techniques developed in this paper are for long-scale fragmentation.

1 Introduction

A major impediment to object recognition is occlusion, in which nearby objects partially obscure distant ones. The bounding contours of the partially occluded objects become fragmented (Fig. 1), forcing us to invent procedures to "link" the fragment endpoints together to form groups of contours belonging to the same object.

The long-distance contour fragment grouping problem is intimately related to the perception of occluders. In Fig. 2, only with the perception of the occluders can the contour fragments be grouped to allow the recoginiton of the Bs. We generalize from this example to suggest that long-distance contour fragment grouping is a kind of inverse to physical occlusion:

Principle of Perceptual Occlusion FORWARD: *Given a collection of objects distributed in space, occlusion will arise generically under projection onto images. The result is a composite bounding contour which, by transversality [18, 24], will contain discontinuities at T-junctions. Consequently, the bounding contours of occluded objects may be* FRAGMENTED *in the image.* INVERSE: *Suppose a collection of contour fragments is given. Long-distance contour fragment grouping is the association of fragments that could have been the result of the occlusion of a single object by some occluder.*

Figure 2: (LEFT) Fragment grouping and occluder description are related. Nine scattered islands of figure fragments with no apparent structure. The mere removal of certain curves (CENTER LEFT) releases the endpoints for interaction [23, 35], leading to the perception of three identical and familiar objects behind illusory white patches. (CENTER RIGHT) Contour fragments only. (RIGHT) Three objects behind solid occluders. Note the similarity to the illusory patches. After Bregman [5].

However, since many occluders might equally account for an observed fragmentation, this inverse problem is ill-posed [17], and so we will prefer simple to complex shapes for the occluder. Thus a pair of fragment endpoints should be linked only if some simple occluder could have produced them. Note that this *shared, simple occlusion* (SSO) constraint is not satisfied for arbitrary pairs of endpoints (Fig. 3), either due to their relative position and orientation or to some other intervening fragment. Our main goal in this paper is to develop these intuitions into a computable constraint to be applied in long-distance contour fragment grouping.

This view is special in three ways. First, it emphasizes the distinction between long and short distance grouping: occluders create large gaps, while sensor noise, for example, creates only slight contour interruptions (Fig. 1). There is no reason to believe that solutions to the extreme form of either problem should be similar [60]. Techniques for short-distance grouping can be based on differential geometry, where the osculating circle, e.g., defines an (infinitesimal) scale over which co-circularity holds [39, 16]. Second, it suggests that fragment grouping only requires weak information—*that* two endpoints are linked, not *how* they are linked in terms of some explicit reconstruction of the occluded contour. This implies a *functional linkage* [41], which we shall formally model as an equivalence class

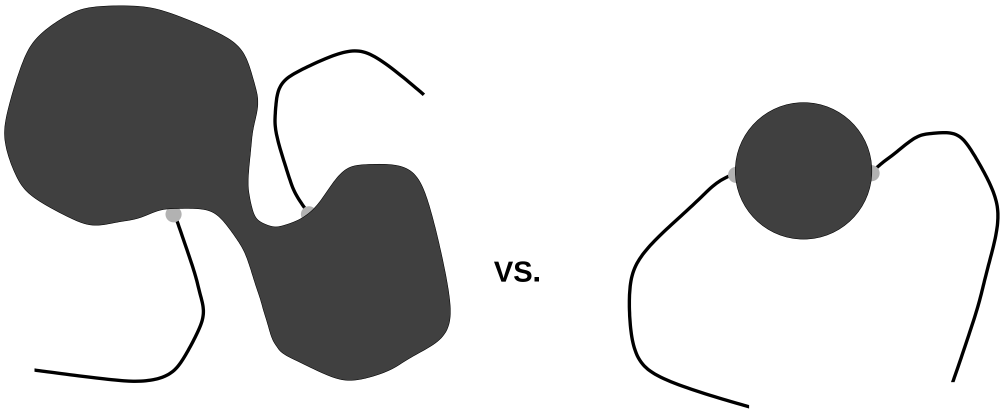

Figure 3: Only a strange, complicated occluder could have caused the occlusion on the left, while a simple object could co-occlude some endpoints on the right.

of possible missing contours. Finally, this approach amounts to hypothesizing a virtual occluder, with both boundary *and* interior region. We shall articulate our notion of occluder simplicity in terms of a classical shape representation—Blum's skeleton—by unifying a subset of the skeleton with endpoints. Together this will define the *gap skeleton*, which includes a description of an implicit occluder (the gap disc) and explicitly indicates which fragments have a shared, simple occluder.

The organization of this paper is as follows. After describing other methods for grouping, we review the skeleton and introduce the gap skeleton as an explicit description of the implicit virtual occluder. We then show how to compute the gap skeleton using curve evolution, present example computations, and then describe virtual occluders. Finally, we demonstrate that the gap disc can be interpreted as an equivalence class of missing contour interpolations using stochastic completion fields [55]. Technical properties related to gap skeleton are proved in an appendix.

2 Background

Most other work on fragment grouping is based on various assumptions about the shape of the occluded object. Contour fragment grouping principles are classified based on whether the information that they exploit derives from endpoints, or from the remaining, non-endpoint, "interior" part of the curves. We shall review the latter first. Before entering the review, however, we stress that this is not a paper on the psychophysics of grouping. Rather, we simply use some of the famous Gestalt examples, and their modern counterparts, as motivation.

2.1 Grouping via Global Shape Assumptions on Partially Occluded Objects

Grouping approaches that use curve interior information often require that the occluded objects have a strict, globally invariant property. For example, those fragments consistent with the projection of a bilaterally symmetric contour can be grouped [32, 58]. This and other classical Gestalt [27] grouping cues, such as parallelism [32] and convexity [22], have been used in perceptual organization systems, in part because psychophysical testing has shown these cues to facilitate shape discrimination and identification [13]. Unfortunately, such global shape assumptions are unnatural: good results can be obtained in an image of disposable razor blades [32], whereas for general smooth blobs the results are less satisfactory. Generalizations to curved symmetry [33], or differential and algebraic invariants [59], still impose serious global restrictions on object shape; in particular, fragments arising from an object which does not satisfy the invariant will not be grouped correctly. In contrast, our shared, simple occlusion (SSO) constraint imposes much weaker assumptions on global shape.

To avoid global shape assumptions, Jacobs instead sought fragment groupings that arise from object parts, since parts are more likely to be convex, say, than the entire object [22]. However, while the parts making up an object may indeed be convex or may come from some limited class, the question of how those elements themselves should be grouped to form a full object representation remains open.

Each of the above constraints holds under specific assumptions, which in turn raises a more troubling question: from a design standpoint, how should the system choose from among the many possible global shape assumptions. In task-specific contexts, perhaps one or several global shape constraints may apply. More commonly, however, arbitrary weightings among the possible constraints have been attempted [42]. It is the intent of the SSO constraint to avoid such choices.

2.2 Grouping via Interpolants: Elastica

Approaches which only use endpoint information attempt to make weaker assumptions about the world by using interpolating curves. For example, elastica are curves that minimize the "energy" $\int(\nu + \mu\kappa(s)^2)ds$ over all smooth curves between the two endpoints, where $\nu, \mu \geq 0$ and κ is curvature [19, 34]. First used to model illusory contours [40, 52], elastica have been the basis for several contour fragment grouping systems [36, 53].

2.2.1 Interlude: Short Distance Assumptions for a Long Distance Problem

Implicit in the use of elastica for grouping is the principle of good continuation. To see this, we note that estimates of elastica are based entirely on information from the fragment endpoints, including derivative information such as direction and curvature, and so they implicitly assume that the entire missing contour is smooth, analogous to Taylor series approximations. When the distance between endpoints is small, this smoothness assumption—the principle of good continuation—is reasonable and is analogous to noise removal using a low-pass filter.

Curve detection [39] is a successful application of good continuation to a grouping problem. Algorithms have been designed that bridge gaps between edge points, but only small gaps caused by sensor noise, for instance, can be handled well [9, 45, 30]. While curve detection is a grouping problem, the elements it operates *upon* are edge points or tangents [2, 4, 10, 14], not the contour fragments that are its *output*. Contour fragment grouping, we maintain, is a distinct and later process, operating over much longer distances such as those that are caused by occlusion. Moreover, given these differences, we cannot support the assumption of a common "grouping engine" either [48].

Global assumptions on the occluded contour Good continuity imposes a global smoothness constraint over the entire occluded contour. Unfortunately, *the smoothness of object boundaries is at best piecewise*, so as the distance between fragment endpoints is increased, it becomes increasingly plausible that the global smoothness constraint will

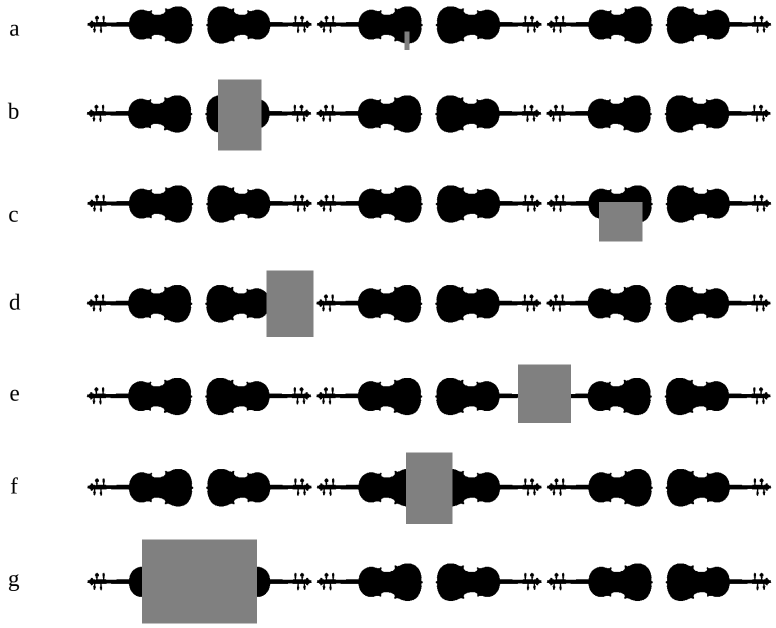

Figure 4: This figure, inspired by Kanizsa [23], shows how the "model" for the occluded contour behind the occluder may switch. Although for a tiny occluder the illusory and occluded contours are similar (a), as the size of the occluder grows, this similarity disappears. For example, the detailed sides of the violin are lost (b), despite strong contextual information, making it appear like a mandolin. In addition, a bilateral symmetry is broken (c), an occluded fingerboard (d) does not perceptually exist, and an occluded pair of violins becomes joined at their necks (e), or elsewhere (f and g). In general, the probability of such a switch should increase with the size of the occluder. Our method (§6) will correctly link the endpoints in all of these examples, and yet does not pick a "model" for the missing contours.

break down at some discontinuity within the missing contour. At long distances, the model appropriate for the contour near the endpoint could be very different from that for contour segments far away: the best model may "switch". Thus endpoint information such as orientation can be misleading, giving rise to interpolants that cannot even qualitatively approximate the missing contour between distant endpoints, such as in the front leg of the camel in Fig. 5. Fig. 4 shows how the interesting details of a violin are not represented visually using any interpolant. Although this is not a perceptually-oriented paper, we note that this breakdown of interpolant models of the occluded contour has been measured perceptually: illusory contours (perceptual interpolants) are not perceived by human observers when

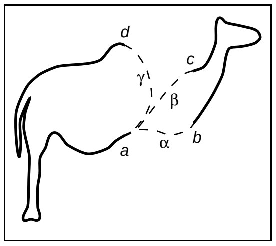

Figure 5: Ambiguity in interpolating between fragments using elastica. "Optimal" continuations α, β, and γ, joining end point a to ends b, c, and d, respectively, all have similar energies based on length and total squared curvature. Note the lack of resemblance to the acutal missing contour.

the inducing endpoints are farther apart than a visual angle of 0.5° [47]. Clearly, contour fragment grouping mechanisms which operate at long distances must have a different basis.

2.2.2 Global Computations and Topological Constraints

Despite these difficulties, elastica have been applied to contour fragment grouping. Indeed, a preference for elastica with low energy (short and relatively straight interpolations) can induce reasonable linkings in some situations. In general, however, the energies of several elastica from a given endpoint may be similar, making such energy-based linkings ambiguous (Fig. 5). Even when combined with other assumptions such as contrast or texture continuity, an optimization computation seeking a globally best selection among all the elastica is required. Since a priori all *pairs* of endpoints are candidates for linking with elastica, this global computation will be combinatorially complex; to cope, arbitrary thresholds on elastica energy have been introduced—a constraint, to be sure, but one whose motivation is less clear than SSO. In addition, these global procedures are typically topologically constrained to ensure that the elastica complete descriptions of surfaces that make physical sense as smooth manifold-solids [57] or paper cutouts [36].

Closure The paper cutout constraint is equivalent to the requirement that the bounding contour of every object is a closed curve. That figural closure is connected to the grouping problem is revealed in Fig. 6, where the three bow-ties to the right are usually more salient than the three double-headed split arrows to the left. While parallelism and proximity might

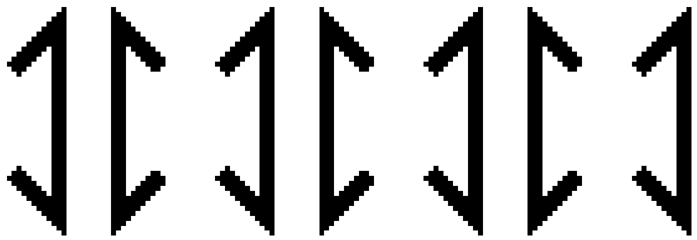

Figure 6: Köhler's illusion. The three bow-ties to the right are seen instead of the pairs of arrows to the left. This is a classical demonstration that closure is involved in fragment grouping.

predict the opposite result, the perceived bows can be explained by a process that extracts groups as sets of fragments that form a *closed* curve. Grouping based on such figural closure is important psychophysically [12]. Closure computations [8, 14, 44] attempt to force all fragments to be elements of closed chains, or cycles ("closed" groups with respect to their endpoints). While closure is sufficient for linking the fragments, closure is not necessary, as some endpoints should never be linked, such as at cusps in the contour generator [26, 467 ff.]. When the three bow-ties form, one fragment is left over and is possibly half of another bow-tie with its endpoints unlinked. Thus closure itself is not always a desirable goal, and its precise role in global grouping computations remains open.

It is clear that some sort of constraint to limit the vast array of grouping possibilities is necessary, and we seek a weaker one than good continuation, based on SSO. This view of grouping is emerging in psychophysics for different reasons, in particular to explain visual search experiments showing no evidence of visual filling-in of missing contours and regions (Rensink and Enns [41]), unlike what grouping theories based on good continuation and illusory contours [40] would suggest. Instead, they find that there is a "functional" filling-in of the missing information, i.e., an abstract association of the literal visual information, in this case, curve fragments. We claim that an implicit occluder both acts as this functional linkage and provides the sought-after constraint.

3 Overview of approach

The minimal universe in which contour fragmentation arises from occlusion is that of contrasting paper cutouts arranged with a partial ordering in depth. This 2.1-D world [36] gives rise to T-junctions when objects occlude, and we shall assume, for the bulk of this paper, that we are in a subclass of this world in which there are only two layers of objects (see §9 for a generalization). We begin with a decomposition at T-junctions as shown for the pair of objects in Fig. 7 (such T-junctions could be found as multiple strong and distinct oriented operator responses at an image position, for example). Thus fragments explicitly arise, and our computational goal is to recover a linking of such fragment endpoints. We seek the groups that most naturally account for the data in the sense of the above principle of perceptual occlusion. As this simple example illustrates, the linking is dictated by the gap skeleton, a certain portion of skeleton related to endpoints (shown in bold, bottom left, Fig. 7). In particular, two curve endpoints are grouped only if they have a gap skeleton. This also provides a rough (skeletal) approximation to an object that, if it were present, would have broken the rectangle into the fragments shown. Such virtual occluders are described in detail in §7 and are shown to cover interpolant estimates of occluded contours with high probability in §8.1. We now formally introduce the gap skeleton.

4 Skeleton and gap skeleton

To formalize the notion of the set of curve fragments to be grouped, we shall consider a closed set $A \subset \mathbb{R}^2$ that is a disjoint union of traces of a finite number of simple, piecewise smooth curves, either open or closed. Recall the definition of the skeleton [3] of a closed set $A \subset \mathbb{R}^2$ in terms of MAXIMAL OPEN DISCS, where an open disc D is maximal if and only if there exists no other open disc contained in the complement of A that properly contains D:

Definition 1 (Blum) *The* SKELETON *S of A is the set of centers of maximal open discs contained in the complement of A.*

Note that the skeleton, which itself is composed of curves, is being used here to describe a set of curves (having empty interiour); typically the skeleton has been applied to the

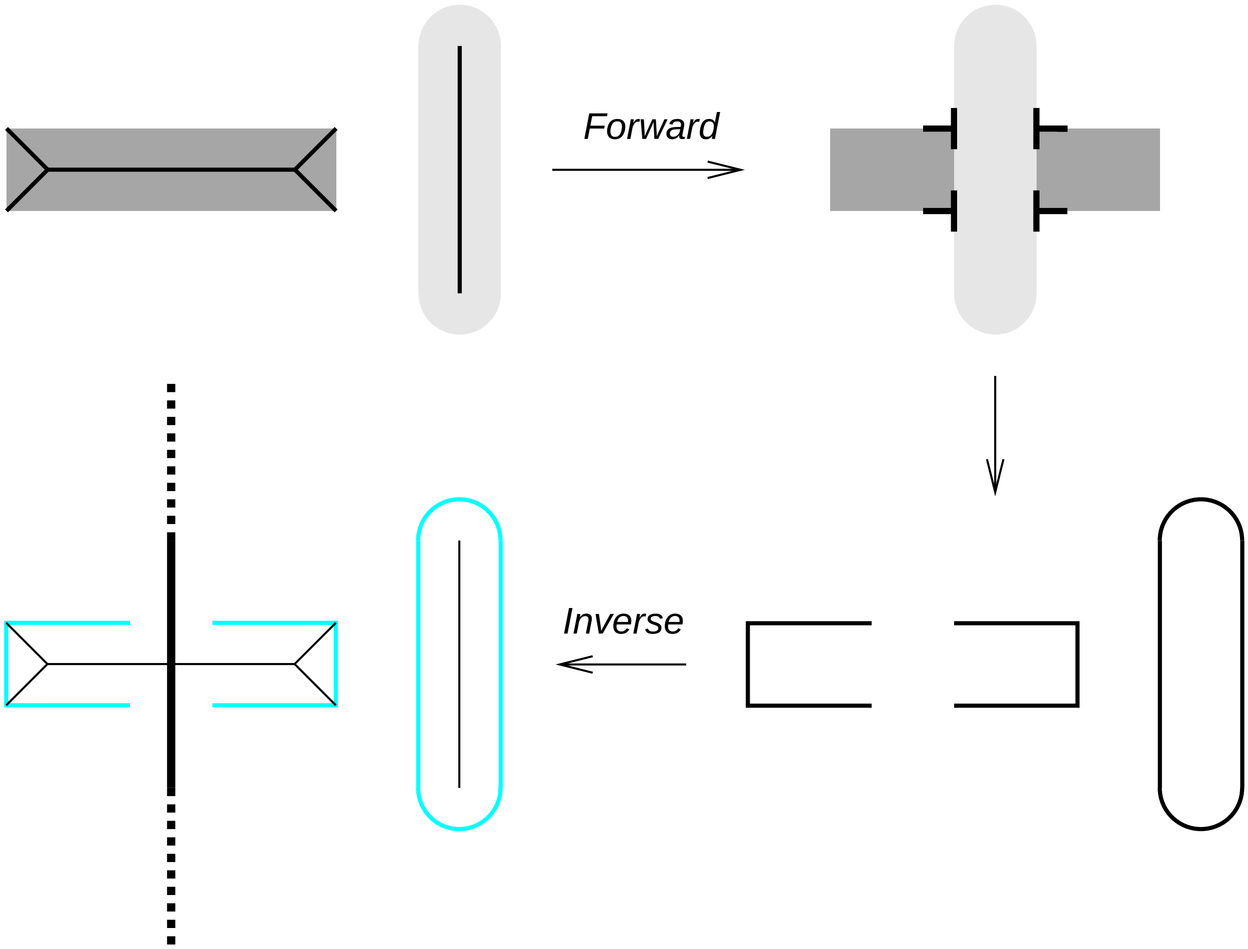

Figure 7: Grouping using the gap skeleton. (TOP LEFT) Two objects with their skeletons. (TOP RIGHT) The two objects in an occlusion relationship; observe the formation of T-junctions. (BOTTOM RIGHT) The bounding contours are separated by a partial ordering in depth induced by T-junctions: the fragments of the rectangle are "farther" than the bounding contour of the occluder. (BOTTOM LEFT) The skeletons of the fragments are shown as thin lines, while the gap skeleton is the thick vertical line. Note the similarity between the gap skeleton and the skeleton of the occluder; consider the gap skeleton as representing a "virtual occluder".

description of shapes (sets with nonempty interiour).

In their study of the skeleton, Calabi and Hartnett introduced the notion of the projection of a point to the nearest point(s) in A [7]. The distance $\rho(x)$ from the point $x \in \mathbb{R}^2$ to the set A is the minimum distance from x to any point of A, or $\rho(x) \triangleq \min\{\|x-p\| : p \in A\}$. The PROJECTION $\pi(x)$ is the set of points in A closest to x, or $\pi(x) \triangleq \{p \in A : \|x-p\| = \rho(x)\}$. The RADIUS FUNCTION r is the restriction $r \triangleq \rho|_S$ of ρ to the skeleton S and so $r(q)$ is the the radius of the maximal disc at q. The projection $\pi(q)$ to A is the set of points at which the maximal disc at q "touches" A. In fact, the intersection of the boundary of the maximal disc at q and A is precisely $\pi(q)$ (see lemma 5 in the appendix). Except at a finite number of points of S, the maximal disc will touch A at two points.

We introduce the gap skeleton through Blum's more general notion of full ligature [3][1]:

Definition 2 *The* FULL LIGATURE *with respect to endpoints a and b is the set* $\mathrm{FL}(a,b) \triangleq \{q \in S : \pi(q) = \{a,b\}\}$.

In other words, $\mathrm{FL}(a,b)$ is that subset of the skeleton whose maximal discs touch a and b (Fig. 8). Note that $\mathrm{FL}(a,b)$ may be empty (e.g., Fig. 3 (left)) and $\mathrm{FL}(b,a) = \mathrm{FL}(a,b)$. The proofs of the following are in the appendix.

Proposition 1 $\mathrm{FL}(a,b)$ *is contained in the perpendicular bisector of the line segment joining a and b.*

Proposition 2 $\mathrm{FL}(a,b)$ *is connected.*

Let θ_a be the signed angle *from* the tangent extending out of endpoint a *to* the line segment $\overline{ab}$. Define θ_b similarly (see Fig. 8).

Proposition 3 *If* $\mathrm{FL}(a,b)$ *is nonempty and* $|\theta_a + \theta_b|$ *is not equal to* $180°$*, then* $\mathrm{FL}(a,b)$ *has nonzero length.*

Coupled with these properties, we see that "almost always" a nonempty full ligature is a "chunk" of the perpendicular bisector between its related endpoints. The addition of a further constraint leads us from full ligature to gap skeleton.

[1] Blum's use of ligature was in terms of concave corners within a closed shape, not the curve endpoints considered here.

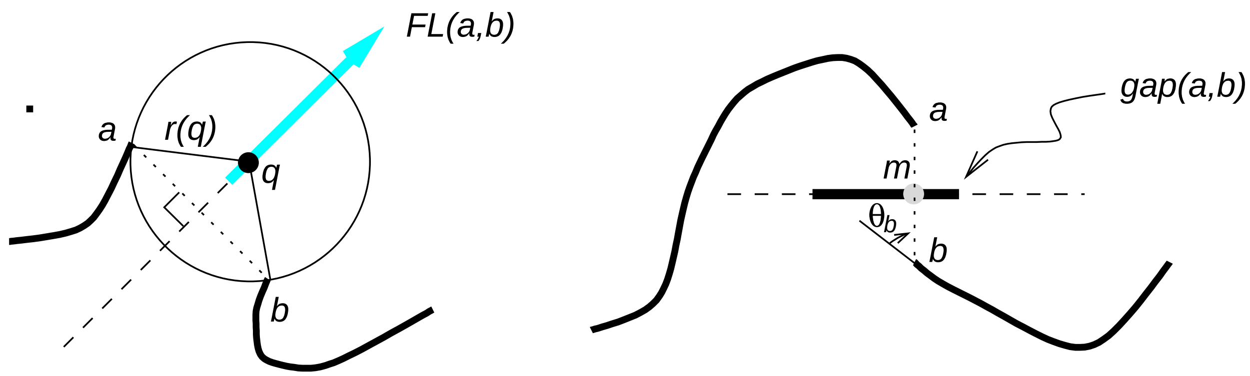

Figure 8: (LEFT) The definition of full ligature is illustrated: only endpoints a and b are a distance $r(q)$ away from the skeletal point q, so $\pi(q) = \{a, b\}$. All such skeletal points form $\mathrm{FL}(a, b)$, and extend along a ray on the perpendicular bisector of the endpoints. Note here that the gap skeleton with respect to a and b is empty because the midpoint is not in $\mathrm{FL}(a, b)$. (RIGHT) Here there is nonempty gap skeleton with respect to a and b, and the radius function r obtains a local minimum at m, the midpoint between a and b.

Definition 3 $\mathrm{FL}(a, b)$ *is called the* GAP SKELETON *with respect to a and b if and only if $r(q)$ achieves a local minimum in* $\mathrm{FL}(a, b)$*:*

$$\mathrm{gap}(a, b) \triangleq \begin{cases} \mathrm{FL}(a, b), & \textit{if } r \textit{ has a local minimum in } \mathrm{FL}(a, b); \\ \varnothing, & \textit{otherwise.} \end{cases}$$

Proposition 4 *The midpoint between a and b lies in* $\mathrm{FL}(a, b)$ *if and only if r has a local minimum in* $\mathrm{FL}(a, b)$*.*

Intuitively, gap skeleton arises when maximal discs touch a given pair of endpoints and become smallest "between" those endpoints. Thus gap skeleton with respect to a pair of endpoints will be present when no curve enters into the *gap disc*, i.e., the disc whose diameter is the line segment joining the endpoints. Note that, for a given pair of endpoints, the presence of full ligature does not imply the presence of gap skeleton, as shown in Fig. 8. In addition, if $\mathrm{FL}(a, b)$ is empty for all pairs of endpoints (a, b) of a given set of curves, then so is the related gap skeleton.

The formalization of the principle of perceptual occlusion is now at hand. We say that *a pair of endpoints satisfies the* SSO *constraint if and only if there is gap skeleton with respect* *to that pair of endpoints*. Operationally, we shall link a pair of endpoints only if they are related via gap skeleton. The gap disc becomes the simple occluder that is shared by a pair of endpoints, and is thus an explicit representation of an implicit occluder (see §7) that we sought in §1.

Interestingly, others have used the related notion of Voronoi tessellation for grouping [1, 20], but in these approaches no special status is given to endpoints. Endpoints and the discontinuities in boundary orientation that give rise to them are critical for unit formation in human psychophysics [24]. Recently, Tek, Stoll, and Kimia have addressed the problem of perceptual organization using the skeleton induced by fragmented curves [51]. In the context of analysing MRI images, Lohmann extracts closed contours using maximal discs as well [31], but requires that an initial seed point be given.

5 Skeletons and shocks

In Blum's *grass-fire* formulation, the skeleton is obtained as quench points of a fire front moving parallel to the shape's boundary. An evolutionary approach to shape description formalizes this view by introducing a partial differential equation [25, 6]: $\frac{\partial C}{\partial t} = N$, with the initial condition $C(s,0) = C_0(s)$, where $C(\cdot,t)$ is the closed curve describing the boundary of the shape after a deformation by time t, N is the outward normal, s is the arc-length parameter. Curves evolving according to the above equation develop *shocks* [28], or entropy-satisfying singularities, which are classified into four types [25], two of which directly relate to the gap skeleton (see Fig. 9). The numerical algorithm that we use to detect shocks [49] provides accurate estimates of the radius function as well as shock velocities, both of which are are critical for the computation of the projection $\pi(q)$ of a shock point q; as a bonus the detected shocks are automatically grouped into shock branches.

As we show, first- and second-order shocks are extremely important. First, however, we note that the curve evolution approach is designed for a collection of *closed* curves in the plane having a topologically well-defined inside (the computation being carried out on an associated embedding surface [38]). Open curves, however, are slightly more subtle since they arise from edges for which the orientation is known modulo 180°: the polarity of the

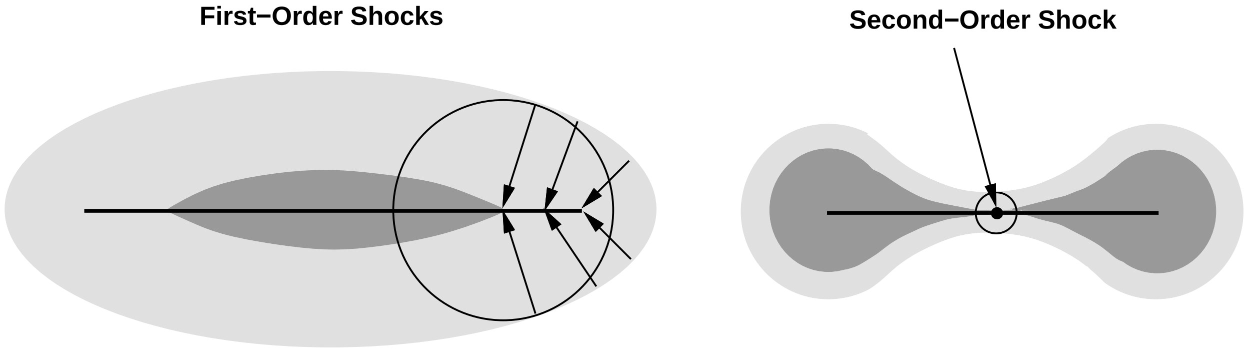

Figure 9: Some shock types which occur during curve evolution. At two shock points the related maximal disc is shown. (LEFT) A FIRST-ORDER SHOCK is a discontinuity in orientation of the boundary of a shape. (RIGHT) A SECOND-ORDER SHOCK is formed when two distinct non-neighboring boundary points join, but none of their immediate neighbors collapse together.

orientation is undefined (see Figures 1 and 7). Thus both polarities must be considered, which we effect by dilating the curve an ϵ-amount. The boundary of the dilation is a closed curve with interior, and evolution can now proceed as usual (see Fig. 10). This informal argument is supported by the topological observation that a curve is the boundary of a set with empty interior; the ϵ-dilation realizes this connection in discrete domains.

We are now in a position to see the relationship between the shocks formed during curve evolution and the definition of gap skeleton. The locus of points through which the shocks migrate corresponds to a Blum skeleton. In Fig. 10 [12], note how the endpoints propagate outwards as semicircular arcs. At the top of the figure, the collision of the two arcs midway between the endpoints causes a second-order shock, which instantly becomes two high-velocity first-order shocks heading in opposite directions. Since the time of shock formation is equivalent to the radius of the maximal disc at the shock (Fig. 9), the gap skeleton of a pair of endpoints is the set of shocks of the corresponding semicircular arcs, provided that this set includes a second-order shock. In summary, we have:

Proposition 5 *A second-order shock q satisfies $\pi(q) = \{a, b\}$ if and only if $\mathrm{gap}(a, b) \neq \varnothing$, where a and b are endpoints.*

We note that Tari, Shah, and Pien have demonstrated an approximation to the grass-fire

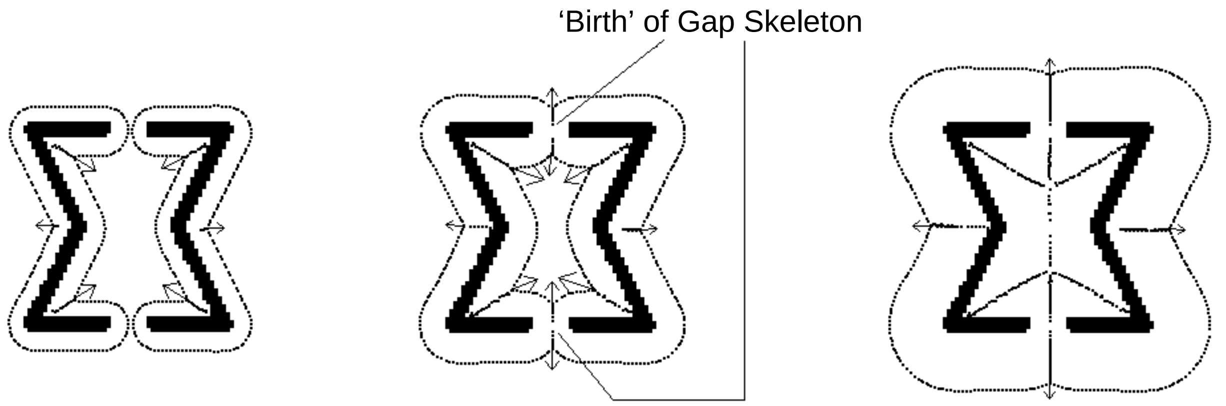

Figure 10: (LEFT) The bold figure is the dilation of the broken bounding contour of an hourglass. A front during curve evolution is overlaid, with first-order shock points forming in high curvature regions. Arrows denote the direction and speed of first-order shocks. Endpoints propagate as semi-circular arcs. (CENTER) Gap skeleton is "born" at both the top and bottom: a second-order shock exists instantaneously, and then two first-order shock branches emerge from this collision of fronts from endpoints. (RIGHT) The front continues to propagate outwards, with first-order shocks tracing out the skeleton.

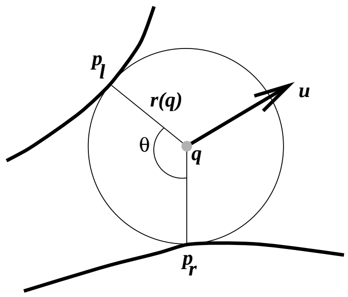

Figure 11: The geometry of the computation of the projection $\pi(q) = \{p_l, p_r\}$ of skeletal point q. The unit vector u points in the direction of the first-order shock.

transform on fragmented figures that exhibited several branches of an approximate skeleton similar to gap skeleton [50]. Having related shocks to the gap skeleton, we now proceed to demonstrate the detection of gap skeleton for real examples.

6 Gap skeleton computations

To compute the gap skeleton, we first compute the projection $\pi(q)$ for all skeletal points q, then detect the full ligature for pairs of endpoints, and finally search for gap skeleton, if any, within the full ligature.

To compute the projection of q, we need not only the skeleton, but also local derivative properties. Specifically, the formulas for the two points, p_l and p_r, in the projection $\pi(q)$ are as follows:

$$\begin{aligned} p_l &\triangleq q - r(q)\mathrm{Rot}(-\theta/2)u, \\ p_r &\triangleq q - r(q)\mathrm{Rot}(+\theta/2)u, \end{aligned}$$

where $\theta = \arcsin(1/s)$, s is the speed of the first-order shock, u is the unit vector in the direction of the shock, and $\mathrm{Rot}(\phi)$ is the two-dimensional rotation matrix through angle ϕ (see Fig. 11) [46]. Thus we require the position, radius, direction, and speed of a skeletal point to compute its projection. Accurate estimates of all these quantities are provided by the algorithm for shock detection [49].

Our method for detecting full ligature points exploits the uncertainty in the detection of endpoints of a curve. Given an endpoint uncertainty ϵ, the full ligature with respect to

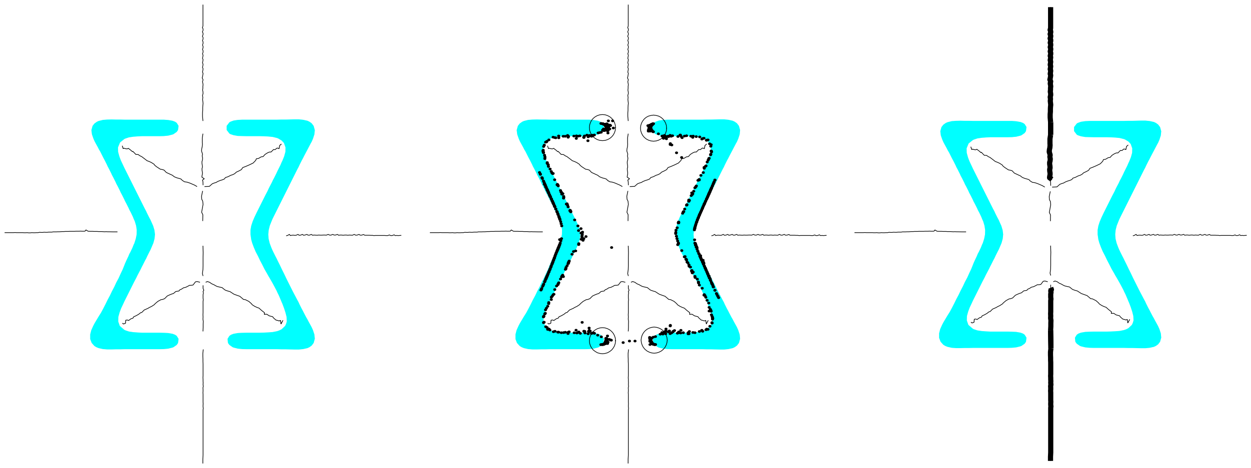

Figure 12: The process of gap detection. (LEFT) The initial level set is displayed in grey and the thin lines are first-order shock branches obtained from its outward evolution. Second-order shocks are not shown, giving a broken appearance to the skeleton. (MIDDLE) The ϵ-balls (circles) around the curve fragment endpoints (manually positioned) and the projection $\pi(q)$ (dots) of each skeletal point q are overlaid. (RIGHT) The detected gap skeleton (thickened lines) is shown; observe that it is connected in accordance to Proposition 2.

endpoints a and b is the set of skeletal points whose projection falls within an ϵ-ball around these endpoints. The connectedness of the one-shock branches from the shock detection and the manner in which one-shock branches flow out of a two-shock branch greatly simplify the computation of the gap skeleton. Essentially we need only search for two full ligature one-shock branches traveling in opposite directions, where each branch is induced by the same pair of endpoints.

We now illustrate the computation of gap skeleton with a variety of examples. For numerical as well as theoretical reasons, the curve flow $\frac{\partial C}{\partial t} = N$ is formulated as the level set evolution of an evolving embedding surface [38]. In our simulations, the initial embedding surface is the signed distance function (slightly blurred to combat discretization) of the set of dilated curve fragments. Fig. 12 depicts the process of gap detection. Note that as predicted the gap skeleton lies on the perpendicular bisector of the line joining two endpoints.

Fig. 13 depicts how grouping via gap skeleton can be derived from a real image. Orientation selective logical/linear operators [21] provide initial estimates of the discrete tangent

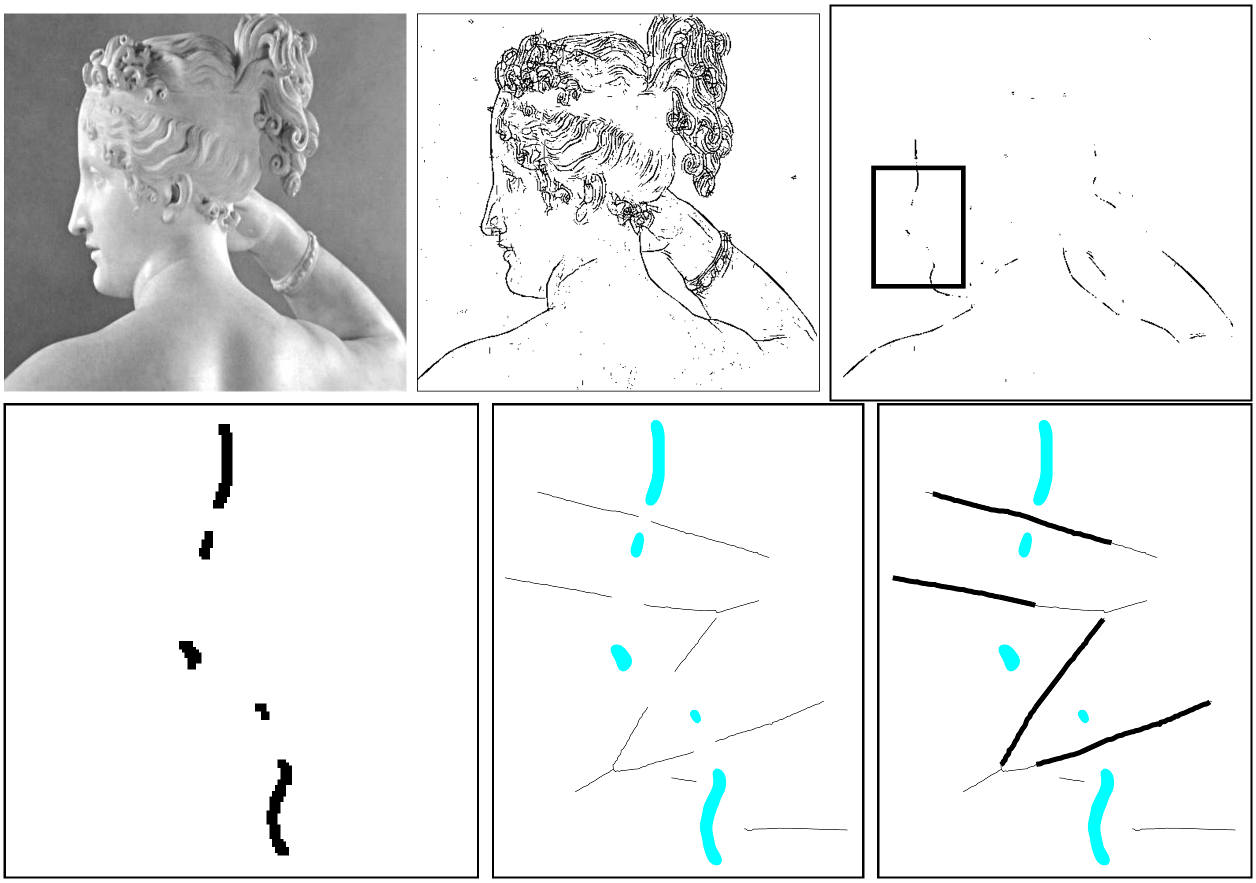

Figure 13: This figure depicts our view for obtaining grouped curves from a real image. (TOP LEFT) Original image. (TOP CENTER) Its discrete tangent map. (TOP RIGHT) We focus on a region of the "contour-like" tangents. (BOTTOM LEFT) The dilated "contour-like" tangents. (BOTTOM CENTER) Shock branches detected under outward evolution are overlaid. (BOTTOM RIGHT) The gap skeleton is shown with thickened lines.

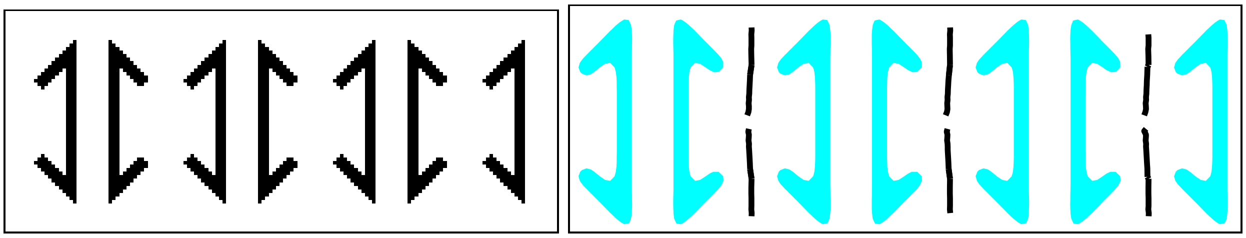

Figure 14: (LEFT) Illusion by Köhler. Note the groupings are seen across the gaps, and not between the parallel sides. Viewers may also see a subjective disc or ellipse (virtual occluder). (RIGHT) The detected gap skeleton (dark lines) selects a grouping of the three "bow-tie" pairs.

map (DTM) of the grey-scale image (Fig. 13, top center), which is then refined using co-circularity [39] implemented by relaxation labeling. A geometrical complexity analysis [11] is then applied to the now geometrically-consistent DTM in order to extract the "contour-like" tangents, which are then dilated (see §5). Connected components in the result are then the contour fragments to be grouped (Fig. 13, bottom left). Shock branches are computed from these curve fragments (Fig. 13, bottom center) and then finally gap skeleton is detected (Fig.13, bottom right). Note that even though the cause of contour fragmentation in this example is not occlusion, the system still works. Fig. 14 illustrates how the gap skeleton explains the bow-tie illusion [27, p. 168, adapted from Köhler]. In Fig. 15, the grouping of the curve fragments from the scene in Fig. 1 is effectively determined by the detected gap skeleton, even though the elastica energies are ambiguous here (see §2).

We note that in general there will be a need to post-process the set of SSO constraint-satisfying endpoint pairs. For example, in Fig. 17, observe at the top-left corner of the top B that one endpoint is related by gap skeleton to two other endpoints. However, even though the SSO constraint does not specify a unique endpoint linking, it clearly reduces the set of possible linkings, and therefore the computational burden of any post-processing (see §9).

7 Virtual occluders

We now return to our motivating principle of perceptual occlusion. Since gap skeleton is a subset of the skeleton, the maximal discs along it can be used to characterize a virtual

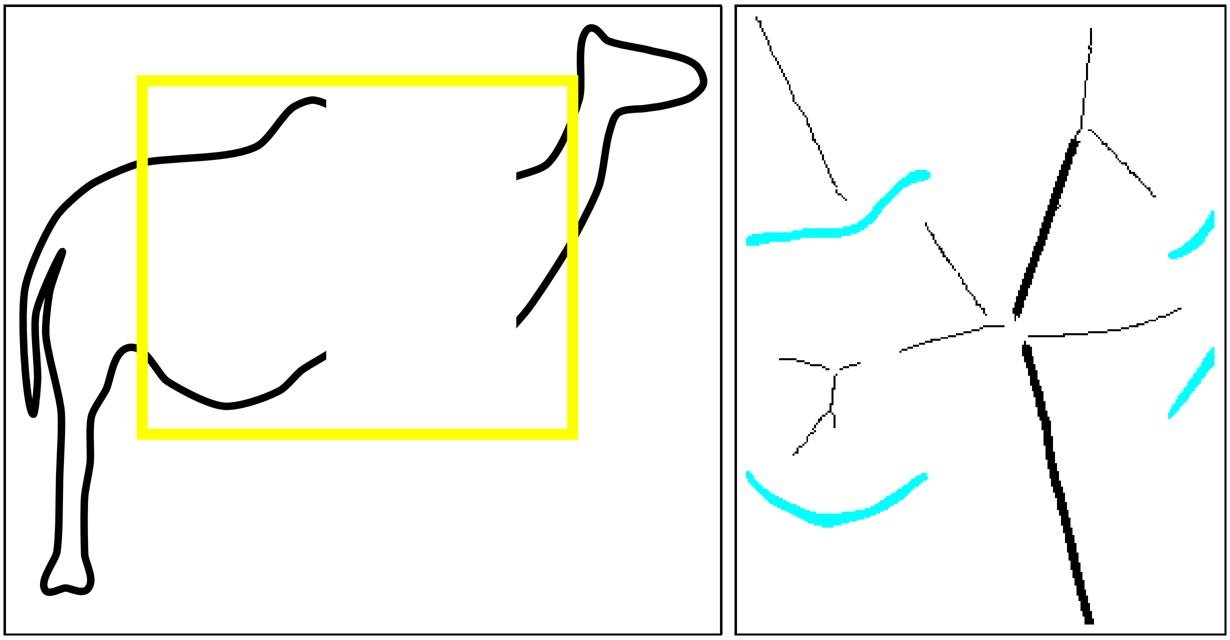

Figure 15: (LEFT) We focus on the region of fragmentation of the camel. (RIGHT) Observe that the gap skeleton (thickened lines) groups the fragments appropriately, appealing to the notion of a virtual occluder; in this case the actual occluder is a tree (see Fig. 1). Notice how the gap skeleton approximates that of the tree.

occluder. A "least commitment" virtual occluder is just the maximal disc at the point along the gap skeleton at which the radius function takes on its local minimum (see Fig. 16, top left). This view in some sense provides the most simple shape for an occluder, for as Blum described [3], a disc is the result of an isotropic growth from an initial point. We call this disc the *gap disc* and view it as a *lower bound* on a virtual occluder.

There is also a sense in which the gap skeleton can constrain a virtual occluder. The maximal disc at a given gap skeleton point q guarantees that there is no contour fragment closer to q than $r(q)$, so the union of maximal discs along a piece of gap skeleton is a region of the plane inside of which no contour fragment lies. If we take the union of such regions for all pieces of gap skeleton, then we arrive at a whole region $\mathcal{U}$ of the plane. We view this region as an *upper bound* on virtual occluders, for the boundaries of the occluders that caused the fragmentations are those which crossed the curve fragments at the endpoints and plausibly lie somewhere in this region $\mathcal{U}$ because $\mathcal{U}$ does not to contain any curve fragments. Fig. 17 strikingly illustrates how the upper bound virtual occluder is present in a fragmented figure and approximately contains the true occluder.

We combine these conceptions of lower and upper bounds in the following construction for virtual occluders. Note that the result is consistent with the constraints in the image:

1. Divide $\mathcal{U}$ into its connected components $\mathcal{U}^1, \mathcal{U}^2, \ldots, \mathcal{U}^n$.

2. For each connected component $\mathcal{U}^i$ generate a virtual occluder as the region inside a closed spline α^i which satisfies the following constraints:

   (a) α^i passes through all of the endpoints bordering on $\mathcal{U}^i$,

   (b) α^i does not cross itself,

   (c) α^i is contained in the upper bound $\mathcal{U}^i$,

   (d) α^i contains all of the lower bound maximal discs that are contained in $\mathcal{U}^i$.

Constraints (a) and (c) attempt to ensure that the missing contours are covered, constraint (d) attempts to ensure that no new occlusions occur, and constraint (b) is necessary topologically. See Fig. 16, bottom. It is an open question how the uniqueness of α^i depends on the choice of spline.

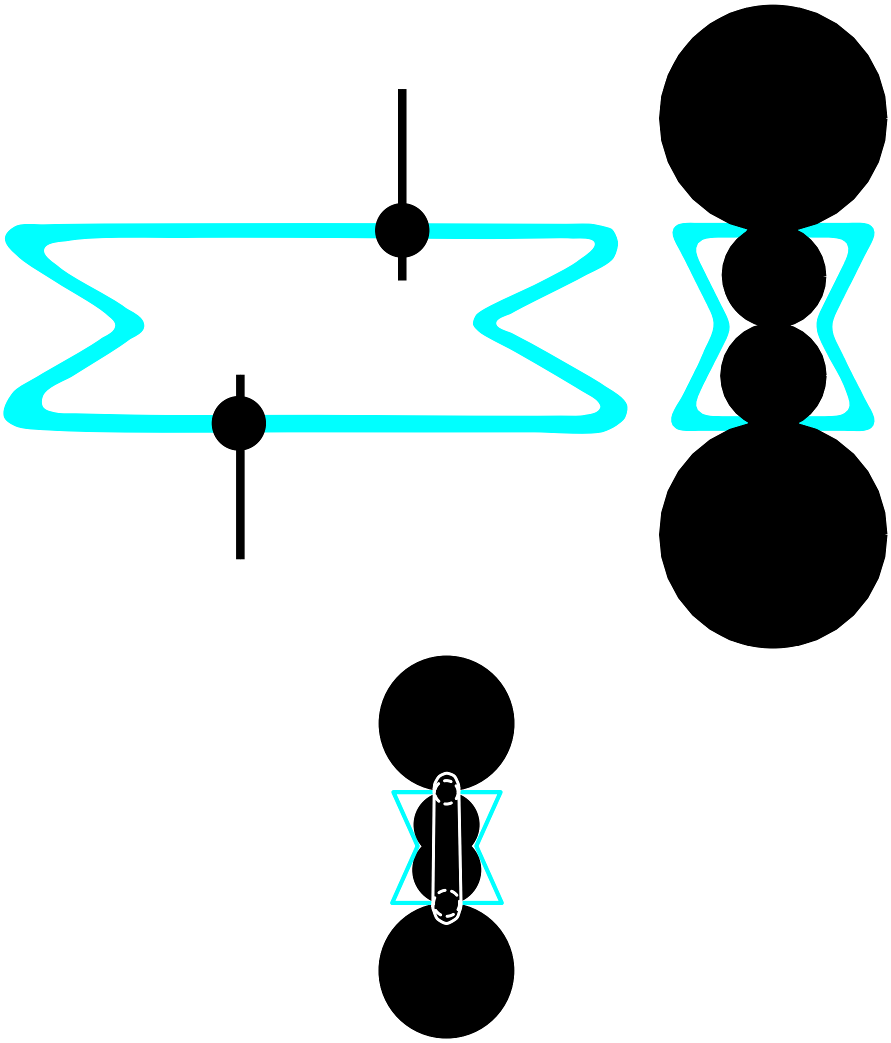

Figure 16: Virtual occluders. Curve fragments are in grey. (TOP LEFT) A disc of radius $r(m)$ (in black) placed where the radius function takes on a local minimum m in the gap skeleton (thick line) is a “least commitment” occluder, or a lower bound on a virtual occluder. (TOP RIGHT) An upper bound on a virtual occluder (black region) is described by a disc at each gap skeletal point q of radius $r(q)$. (BOTTOM) The virtual occluder created using a spline-like curve obeying both the upper and lower bounds (with dashed circles expressing constraint (d)).

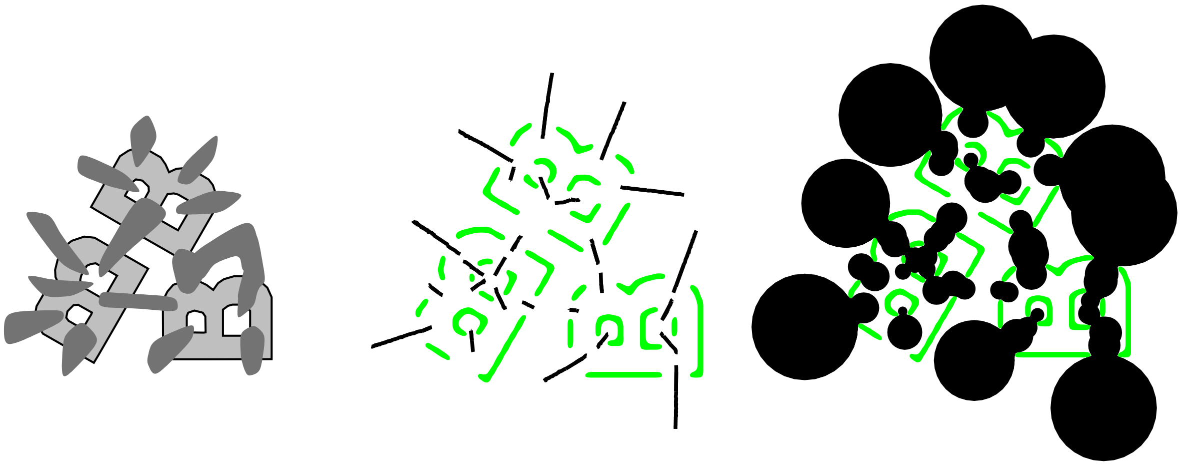

Figure 17: (LEFT) Three objects behind many occluders. (CENTER) Corresponding endpoints are linked *behind* the true occluders by gap skeleton. The two missing linkages are due to discretization (step size of the curve evolution). One endpoint in the top B forms gap skeleton with two endpoints: the SSO constraint is not uniquely satisfied. (RIGHT) The (upper bound) virtual occluders approximately covers most of the true occluders.

There is clearly some degree of flexibility in making implicit virtual occluders explicit. Sometimes they are perceptually unique, while at other times they may not be unique nor even salient. The construction described in this section is sufficient but not perceptually necessary. Additional research on the perceptual viability of virtual occluders is required.

8 Interpretation of the Shared, Simple Occlusion (SSO) Constraint

As the SSO constraint is unlike other conceptions of the fragment grouping problem, it is instructive to clarify its interpretation beyond the intuitive idea of virtual occluders. To do so, in this section we first quantitatively demonstrate how the gap disc covers the occluded contour with high probability. We then show how a particular operationalization of the SSO constraint solves the fragment grouping problem in a simple micro-world—consisting of two layers with disjoint disc occluders—and view its application outside this world as a useful heuristic bolstered by much more heretofore unexploited information.

8.1 Relating the Gap Disc to Interpolating Curves

Whereas the choice of a gap disc may at first seem arbitrary, there is in fact a connection to the missing contour. Since the occluded contour is unknown, we model, following Mumford, Williams, and co-workers [34, 55], the missing curve as a kind of random walk, where a particle travelling tangentially from one curve fragment endpoint has its direction randomly altered as it moves on a path to the endpoint of some other curve fragment. We will demonstrate that, on average, the particle spends almost all of its time behind the gap disc, our simple occluder model. Stated differently, the gap disc covers the overwhelming majority of the possible boundary completions.

8.1.1 Stochastic Completion Fields

Formally, the particle is undergoing a random walk in (x, y, θ) with the stochastic differential equation (SDE):

$$\frac{dx}{dt} = \cos\theta, \ \frac{dy}{dt} = \sin\theta, \ \frac{d\theta}{dt} = N(0, \sigma^2),$$

where (x, y) is the particle's position, θ is its orientation, t is the time in the particle's history, and σ^2 is the variance of the 0-mean white noise process $N(0, \sigma^2)$. The corresponding evolution of the particle's probability density $p(x, y, \theta, t)$ is described by the following Fokker-Plank partial differential equation (PDE):

$$\frac{\partial p}{\partial t} = -\cos\theta \frac{\partial p}{\partial x} - \sin\theta \frac{\partial p}{\partial y} + \frac{\sigma^2}{2}\frac{\partial^2 p}{\partial \theta^2} - \frac{1}{\tau}p,$$

where the last term was added to describe the decay of the particles with a time constant τ.

Now, suppose the missing contour has endpoints (x_1, y_1, θ_1), the "source", and (x_2, y_2, θ_2), the "sink". With the initial condition $p(x, y, \theta, 0; x_1, y_1, \theta_1) \triangleq \delta(x - x_1, y - y_1, \theta - \theta_1, 0)$, we solve the Fokker-Plank PDE[2] we get the source field:

$$P(x, y, \theta; x_1, y_1, \theta_1) := \int_0^\infty p(x, y, \theta, t; x_1, y_1, \theta_1) dt.$$

[2]The Fokker-Plank PDE was solved using the method similar to [56], on a 64-by-64-by-24 grid ($\sigma^2 = 0.01, \tau = 10$) for all examples, except for Fig. 20, where a grid of 32-by-32-by-8 was used to reduce computation time.

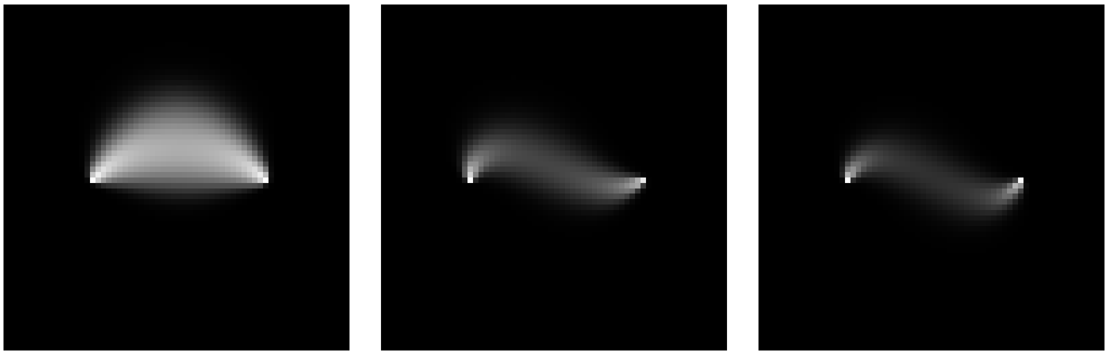

Figure 18: Stochastic completion fields for various endpoint configurations. To display this three-dimensional density on a plane, we integrate over θ.

The source field is the probability density that a particle reaches (x, y, θ), given that it started at (x_1, y_1, θ_1).

The analogous sink field $Q(x, y, \theta; x_2, y_2, \theta_2)$ is the probability density that a particle passes through (x, y, θ), given that it eventually passes through (x_2, y_2, θ_2). The *stochastic completion field* (SCF) is the product of the source and sink fields, or:

$$C(x, y, \theta; x_1, y_1, \theta_1; x_2, y_2, \theta_2) = P(x, y, \theta; x_1, y_1, \theta_1)Q(x, y, \theta; x_2, y_2, \theta_2),$$

and is the probability density that a particle passes through (x, y, θ), given that it started at endpoint 1 and ended at endpoint 2. Several example SCFs are shown[3] in Fig. 18. As proved in [55], the elasticum is the most likely path between two endpoints. These examples show, however, that the SCF is diffuse, and not necessarily thin and curve-like, suggesting that good continuation is a relatively poor indicator of endpoint linking, even under the assumption that the particle undergoes only slight perturbations as it walks (i.e., even if there is no discontinuity in the missing continuation).

8.1.2 The Average Stochastic Completion Field in the Gap Disc

We focus on the situation where occluders are large and therefore endpoint orientation is likely to be misleading. In other words, since it is likely that there exists a model switch within the missing contour, we now deliberately ignore endpoint orientation. Con-

[3]The brightness is proportional to the logarithm of the SCF, following [55].

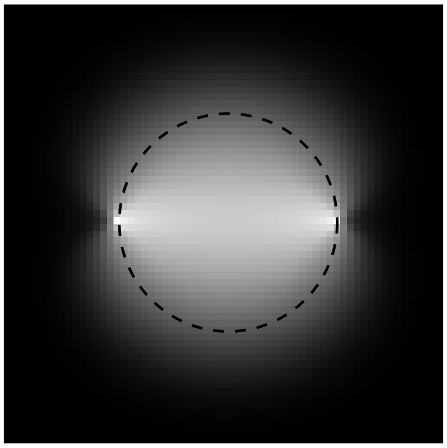

Figure 19: The average stochastic completion field $\overline{C}$. Superimposed is the border of the gap disc, which contains about 98% of $\overline{C}$, for $\sigma = 0.01$ and $\tau = 10$.

sequently, we suppose that the two endpoints are fixed, but allow random[4] endpoint orientations. The average SCF $\overline{C}(x, y, \theta)$ over all endpoint orientations is proportional to $\int_0^{2\pi}\int_0^{2\pi} C(x, y, \theta; x_1, y_1, \theta_1, x_2, y_2, \theta_2) d\theta_1 d\theta_2$ (see Fig. 19). The fraction of the average SCF contained in the gap disc is:

$$\gamma(\sigma^2, \tau) \triangleq \frac{\int_0^{2\pi}\int_{\text{gap disc}} \overline{C}(x, y, \theta) dx dy d\theta}{\int_0^{2\pi}\int_{\mathbb{R}^2} \overline{C}(x, y, \theta) dx dy d\theta}.$$

γ represents the fraction of time that the missing contour remains occluded by the gap disc. We emphasize that γ is especially relevant as the distance between the endpoints becomes large. Observe that γ does not strongly depend on the parameters σ^2 and τ, as is illustrated in Fig. 20, and is typically well above 95%.[5]

Thus we see that, on average, most of the missing contour lies behind the gap disc. Since this least commitment occluder is smallest, the other occluders suggested in §7 will cover even more of the missing continuation. We are encouraged that virtual occluders cover the missing continuation with high probabily, as this supports the notion that the virtual occluder represents an equivalence class of many possible missing continuations. Practically,

[4] Assume orientation is uniformly distributed over $[0, 2\pi)$

[5] To ensure the numerical stability of the Fokker-Plank PDE [56], σ^2 has an upper bound of $2(\Delta\theta)^2$, and is typically set much smaller than this limit [54]. The upper limit on τ is set by the size of pixel grid. The lower limits are 0.

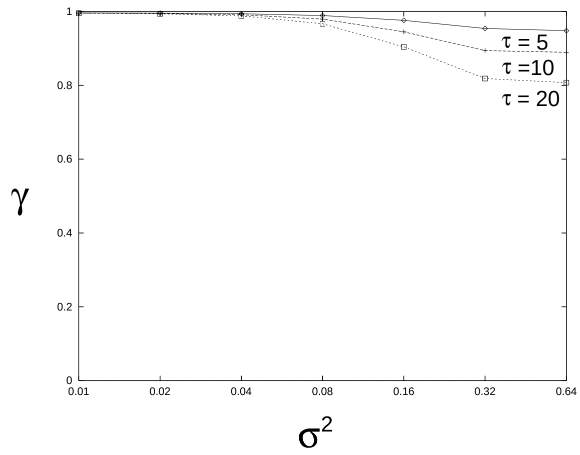

Figure 20: The fraction of the average SCF contained in the gap disc, as a function of the random walk model parameters σ and τ. Typically, $\sigma^2 < 0.05$ [54]. Thus the particle representing the missing contour spends most of its time behind the gap disc.

this suggests that the shared, simple occlusion is a powerful heuristic for constraining which endpoints to link.

8.2 Full Ligature is *Necessary* within a Micro-World

To further articulate the importance of the SSO constraint for contour fragment grouping, we will show in this section that in a significant but limited micro-world, related ideas hold strictly. In particular, we consider those endpoint pairs merely related by full ligature, without the local minimum criterion specified in Def. 3. Observe that a given endpoint will generally be related to more endpoints via full ligature than via gap skeleton, causing a more complex post-processing (see end of §6). We focus on full ligature here instead of gap skeleton because of the special role it plays in this micro-world.

Each scene in our micro-world will be generated by two layers. Suppose that each of the n possibly occluded *objects* is the closure $B_i \subset \mathbb{R}^2$ of an open set with piecewise smooth boundary ∂B_i, where $i = 1, \ldots, n$. Further suppose that the objects are disjoint: $B_i \cap B_j = \varnothing$, if $i \neq j$. Let the bottom layer be defined as the (disjoint) union $B = \cup_{i=1}^{n} B_j$ of objects (Fig. 21, top left). Let the top layer of *occluders* be defined as a (disjoint) union $D = \cup_{j=1}^{m} D_j$ of discs D_j (closed balls $\subset \mathbb{R}^2$), where $j = 1, \ldots, m$, such that the sets B and D meet transversally[6] [18, 24]. Now these two layers are brought into an occlusion relationship with respect to the optical axis of the viewer to produce the scene (Fig. 21), assuming that B and D are constructed by contrasting paper cutouts. As described in §3, the boundaries of those occluding discs that partially occlude the boundary of some object can be removed (due to T-junctions) from the composite set of contours in the scene, and the remaining set of contours, some of which may be open, i.e., due to fragmentation, will be denoted A (cf. §4). Given any pair $\{a, b\}$ of endpoints in A, $\mathrm{FL}(a, b)$ is either empty or not. Is there any guaranteed property of full ligature here?

Suppose object B_i is partially occluded by some disc D_j, and consider any connected subset C_k of $\partial B_i \cap D_j$. C_k is called an *occluded contour* of B_i; the endpoints of any such occluded contour C_k will clearly be located at T-junctions, and thus these will be endpoints of A. The sought-after property of full ligature follows.

Proposition 6 (Full Ligature Necessity for Transversal Occlusion by Disjoint Discs) *If a and b are endpoints of an occluded contour in the above 2-layer world, then* $\mathrm{FL}(a, b)$ *is nonempty.*

PROOF: This claim is true by construction: each disc creates endpoints on the curves it partially occludes. These endpoints abut a disc—the occluder—that crosses no other fragments, and thus the center of this disc will be a point in the skeleton of the fragments. Discs inducing full ligature for the endpoints touching this occluding disc will exist in its neighbourhood as long as we have a generic occlusion, i.e., the occlusion occured transversally. □

Thus we can say that full ligature is a necessary condition for the correct linking of contour fragments in our 2-layer disc-occluder world: all the endpoints that "should" be linked are to be found among the full-ligature-related endpoint pairs. It is interesting to note that discs also play a predominant role in Lee and Mumford's recent model accounting for the scale-invariance of the statistics of natural images [29].

One can show that the condition on the disjointness of the occluders cannot be removed[7]. In addition, the above proposition only claims that full ligature is a necessary condition for

[6]Loosely, transversality in this context means that at any point p where ∂D_j intersects ∂B_i, for some i and j, the tangents to each boundary are distinct.

[7]Consider two nearby narrow sausages, one of which is occluded by two barely-overlapping discs.

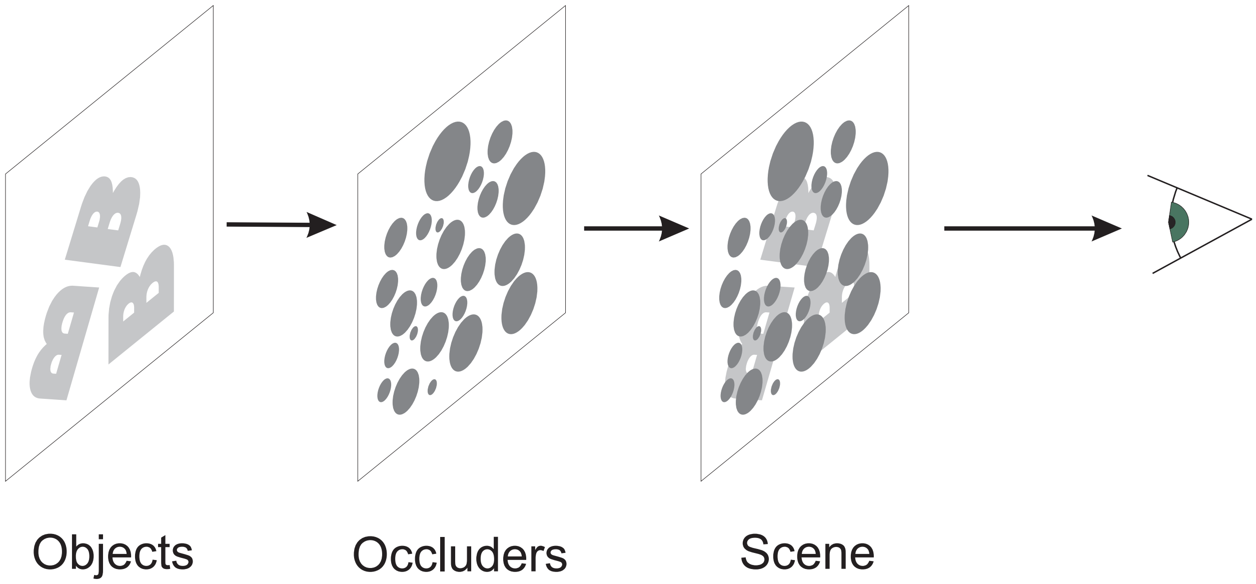

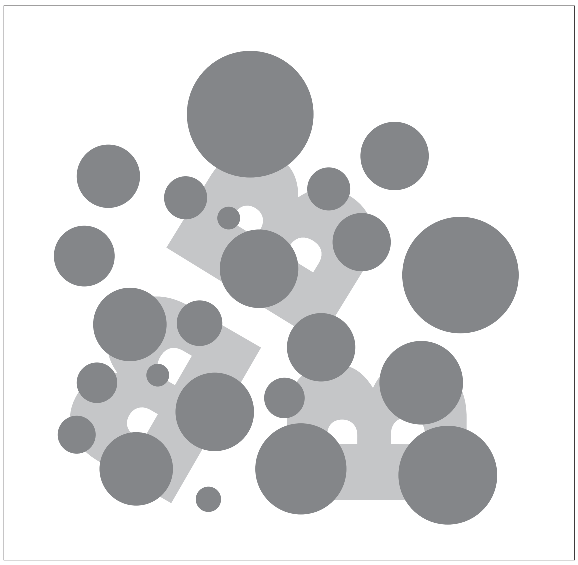

Figure 21: (TOP) A 2-layer micro-world with disc occluders. (BOTTOM) The contour fragments of any scene from this world, such as that depicted here, will be linked correctly only if the endpoints are related by full ligature (Prop. 6).

linking. It definitely does happen that a given endpoint will participate via full ligature with several, and possibly many, other endpoints (gap skeleton is typically satisfied for fewer endpoint pairs). However, even with full ligature, there are at most $O(n)$ pairs of endpoints related by full ligature, a significant reduction over the $O(n^2)$ total possibilities, where n is the number of endpoints[8]. Most other approaches must a priori consider linking with all other endpoints. The computational advantage of the SSO constraint is clear.

Kellman and Shipley's notion of "relatability", in which a pair of endpoints is relatable only if the interpolating curve has no inflection point, is also a constraint for fragment grouping [24], but no worst case improvement over $O(n^2)$ can be made. More importantly, relatability has no explicit place for the occluders or the space they occupy. In addition, relatability imposes a tight constraint on endpoint orientations, excluding many reasonable endpoint linkings (e.g., Fig. 1 and 4): at *long* distances, endpoint orientation is a poor factor for determining linkings. For short-distance fragment grouping, however, where endpoint orientation is more significant, relatability may be useful.

Our 2-layer world is important because it has no constraint on the size of the disc occluders, unlike short-distance good continuation models. While other long-distance grouping approaches exist, such as those based on symmetries or invariants (see §2), this is, to our knowledge, the first generic approach to long-distance fragment grouping.

8.2.1 Full Ligature is Heuristic Outside this Micro-World

Clearly, the visual world is much more general. Say, for example, the occluders were allowed to be of arbitrary shape. It could happen that a pair of endpoints induced by an arbitrary occluder might not be related via gap skeleton: imagine an occluder as an ϵ-dilation of a subset of the boundary of a non-convex object. We conjecture, however, that such failures of the SSO constraint are more pathological than generic, and are encouraged by the examples in §6. In particular, while generally occluders in the world are not discs, they often produce fragmentations which could have been caused by discs, and so we can view the disjoint discs as lower bounds on occluders, in the sense of §7.

[8] To see this, consider generalizing the well-known result [37] that the voronoi diagram for n points is of size $O(n)$

A more difficult issue than occluder shape is occluder layering: suppose there are more than two layers of overlapping objects in the scene. Clearly, the presence of contour fragments from the multiple layers will interfere with the interaction of fragments from a given layer, and gap skeleton will induce inappropriate linkings. However, without more information, humans also have difficulty in performing such groupings (Fig. 22). While this is not a paper about psychology, we conjecture, in contradiction to [24, 53], that a great deal of computation preceeds contour fragment grouping in humans, including the semantic separation of contour fragments. Geiger et al also exploit the layering of fragments [15, §5].

8.2.2 Additional Information Differentiates Contour Fragments

Consider the variety of sources of information that humans have for performing this semantic separation of contours. First, some contour fragments are neither contrast nor texture edges: positive and negative contrast *lines*, perhaps caused by surface markings or highlights (curve-like specularities). Lines have distinct intensity profiles and can therefore be locally distinguished from the edges we are considering [21]. Second, contours terminating at cusps in the surface projection have unique intensity signatures in the vicinity of the endpoints [26]. Third, there is explicit information that induces a layering of the fragments, including stereo and motion discontinuities, colour and contrast differences, and even T-junctions. This surfeit of additional information allows the reduction of the multilayer grouping problem to one within layers. The implementation of this fusion of information is a subject of future research.

9 Conclusion

We have emphasized in this paper how all grouping problems are not similar. In particular, occlusion creates large gaps in curves, creating a long-distance contour fragment grouping problem, in stark contrast to noise related and small contour interruptions. The Principle of Perceptual Organization was introduced to guide us to seek shared, simple occluders (SSO) that could have caused the fragmentation, not estimates of the missing contour. The virtual occluder is an explicit description of this implicit cause of fragmentation.

Figure 22: The effect of layers. (TOP) The grouping of the mess of curve fragments is extremely difficult until the layering structure is made explicit (BOTTOM). Observe how the principle of perceptual occlusion can be invoked to correctly group within each layer. Instead of distinct shading, one could use stereo or motion (kinetic depth) to induce the fragment layers.

But why do we consider virtual occluders when the actual occluders are visible? Although beyond the scope of this paper, we are interested in developing a global grouping algorithm which exploits both virtual occluders and actual occluders (if available) in the following sense: *the virtual occluder should have a description that is consistent with that of the actual occluder.* Thus the virtual occluder concept is useful because it allows us to explicity check whether the principle of perceptual occlusion is satisfied. The virtual occluder is a *hypothesis* provided by a given grouping that can then be *verified* by comparison with the actual occluder. We conjecture that different possibile groupings, although indistinguishable by the SSO constraint, can be contrasted and evaluated using virtual occluders.

A Proofs of full ligature properties

The proofs for the basic properties of full ligature are presented in this section, and, although lengthy, the methods used are elementary, requiring only a basic understanding of analysis at the level of Rudin [43]. Points a and b are endpoints.

Proposition 1 $\mathrm{FL}(a, b)$ *is contained in the perpendicular bisector of the line segment joining a and b.*

PROOF: The proof is by geometric construction. Since the circle of radius $r(q)$ at any $q \in \mathrm{FL}(a, b)$ intersects A at a and b, q is the apex of an isosceles triangle whose base is the line segment joining a and b. The result follows. □

Proposition 2 $\mathrm{FL}(a, b)$ *is connected.*

PROOF: By a change of coordinates involving only a rotation and a translation, we can assume that a and b lie along the ordinate and are both at a distance h from the origin, so that $a = (0, h)$ and $b = (0, -h)$. By proposition 1, $\mathrm{FL}(a, b)$ is contained in the abscissa, and so any two points in $\mathrm{FL}(a, b)$ will be of the form $c = (\tilde{c}, 0)$ and $d = (\tilde{d}, 0)$, where $\tilde{c} < \tilde{d}$ without loss of generality. It is sufficient to show that pre(a,b) is convex, for convex sets are connected [43, pp. 44–45]. Recall that a set X is convex if and only if $\forall x, y \in X, \forall t \in [0, 1], z \triangleq (1 - t)x + ty \in X$. To show convexity we need to show that any point $e = (\tilde{e}, 0)$ is also in $\mathrm{FL}(a, b)$, where $\tilde{c} < \tilde{e} < \tilde{d}$. That is so if e is both a skeletal point and $\pi(e) = \{a, b\}$, by the definition of full ligature.

By lemma 6 it is sufficient to show that $\pi(e) = \{a, b\}$ for e to be a skeletal point. Let $B(x, \epsilon)$ denote the open disc of radius ϵ centred at x. Recall that ∂X denotes the boundary of a set X. Construct an open disc $B(e, \|e - a\|)$ at e and notice that it "touches", i.e., its boundary intersects, points a and b. By showing that $B(e, \|e - a\|)$ is contained in the union of $B(c, r(c))$ and $B(d, r(d))$ (lemma 1), we know that $B(e, \|e - a\|)$ contains no points of A and so we easily get $\|e - a\| = \rho(e)$ (lemma 2). Therefore, a and b are the *only* points of A

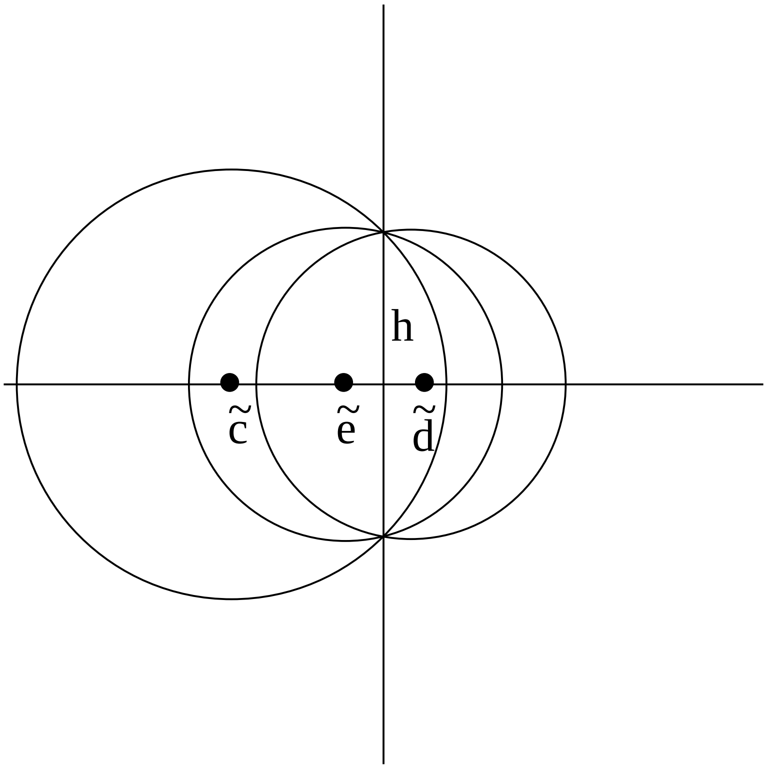

Figure 23: An illustration of lemma 1.

on the boundary $\partial B(e, \rho(e))$ of $B(e, \rho(e))$ because otherwise one of the maximal discs at c and d would contain a point of A (lemma 3). Using lemma 5 we get $\pi(e) = \{a, b\}$, and so $e \in \mathrm{FL}(a, b)$. Thus the desired convexity of the full ligature is shown. $\square$

Lemma 1 $B(e, \|e - a\|) \subset B(c, r(c)) \cup B(d, r(d))$

PROOF: See Fig. 23. Recall that $B(c, r(c)) = B(c, \|c - a\|) = B((\tilde{c}, 0), \sqrt{h^2 + \tilde{c}^2}) = \{(x, y) : (x - \tilde{c})^2 + y^2 < h^2 + \tilde{c}^2\ \}$; $B(d, r(d))$ and $B(e, \|e - a\|)$ can be described similarly. Fix a point $(x, y) \in B(e, \|e - a\|)$, and so $y^2 < h^2 + \tilde{e}^2 - (x - \tilde{e})^2$. A simple calculation shows that $(x - \tilde{c})^2 + y^2 - h^2 - \tilde{c}^2 < 2x(\tilde{e} - \tilde{c})$. Thus if $x \leq 0$, then $(x, y) \in B(c, r(p))$, since $\tilde{e} > \tilde{c}$. Similarly, if $x \geq 0$, then $(x, y) \in B(d, r(d))$. $\square$

Lemma 2 $\|e - a\| = \rho(e)$

PROOF: Since $a \in A$, $\|e - a\| \geq \rho(e)$. Suppose $\|e - a\| > \rho(e)$. Then there exists $p \in A$ such that $\|e - a\| > \|e - p\|$, which is greater than or equal to $\rho(e)$ by definition; hence

$p \in B(e, \|e-a\|)$. So using lemma 1 we know $p \in B(c, r(c)) \cup B(d, r(d))$. This implies that one or both of $B(c, r(c))$ and $B(d, r(d))$ is not a maximal disc, which is a contradiction. Thus $\|e-a\| = \rho(e)$. $\square$

Lemma 3 $\partial B(e, \rho(e)) \cap A = \{a, b\}$

PROOF: By construction, $\partial B(e, \|e-a\|) \supset \{a, b\}$. Since $\|e-a\| = \rho(e)$ (from lemma 2), we have $\partial B(e, \rho(e)) \cap A \supset \{a, b\}$.

We now need to show that $\partial B(e, \rho(e)) \cap A \subset \{a, b\}$. Suppose this is not true; i.e., there exists $g \in \partial B(e, \rho(e)) \cap A$ such that $g \neq a$ and $g \neq b$. By a proof similar to that of lemma 1, we get $\overline{B}(e, \|e-a\|) \subset \overline{B}(c, r(c)) \cup \overline{B}(d, r(d))$, where $\bar{X}$ denotes the closure of X. This implies $\partial B(e, \rho(e)) \subset \overline{B}(c, r(c)) \cup \overline{B}(d, r(d))$, because of lemma 2 and the fact that $\partial X \subset \bar{X}$ for any set X in a topological space. Suppose $g \notin \overline{B}(c, r(c))$. Then $g \in \overline{B}(d, r(d))$. Now $g \notin B(d, r(d))$, for otherwise $B(d, r(d))$ is not a maximal disc because $g \in A$. Thus $g \in \partial B(d, r(d))$ and so the circles $\partial B(e, \rho(e))$ and $\partial B(d, r(d))$ have the points a, b, and g in common. Therefore $\partial B(e, \rho(e)) = \partial B(d, r(d))$, which is a contradiction since $\tilde{e} < \tilde{d}$. Similarly, if we suppose $g \notin \overline{B}(d, r(d))$, then $\partial B(e, \rho(e)) = \partial B(c, r(c))$, which is a contradiction since $\tilde{e} > \tilde{c}$. The result follows. $\square$

Lemma 4 *For all* $x \in \mathbb{R}^2, B(x, \rho(x)) \cap A = \varnothing$.

PROOF: Pick any $y \in B(x, \rho(x))$ and $p \in A$. By the assumptions on A, A is a closed set. Using $\rho(x) \triangleq \min\{\|x-p\| : p \in A\}$, observe that $\|x-y\| < \rho(x) \leq \|x-p\|$. Therefore there do not exist y and p such that $\|x-y\| = \|x-p\|$, which is a necessary condition for $y = p$. Thus $B(x, \rho(x)) \cap A = \varnothing$. $\square$

Here we formalize the idea that a maximal disc touches only the projection of its skeletal point.

Lemma 5 *For all* $x \in \mathbb{R}^2, \pi(x) = \partial B(x, \rho(x)) \cap A$.

PROOF: $\pi(x) = \{p \in A : \|x-p\| = \rho(x)\} = \{p \in \mathbb{R}^2 : \|x-p\| = \rho(x) \text{ and } p \in A\} = \{p \in \mathbb{R}^2 : p \in \partial B(x, \rho(x)) \text{ and } p \in A\} = \partial B(x, \rho(x)) \cap A$. $\square$

Lemma 6 *If the cardinality of $\pi(x)$ is at least two, then x is a skeletal point.*

PROOF: Suppose that the cardinality of $\pi(x)$ is at least two and x is not a skeletal point. Then there exists $y \neq x$ such that $B(y, \rho(y)) \supsetneq B(x, \rho(x))$, and so $\rho(y) > \rho(x)$. Define a and b through $\{a, b\} \subset \pi(x)$, where $a \neq b$. We shall show that at least one of a and b is in $B(y, \rho(y))$. Suppose $a \notin B(y, \rho(y))$. Since $a \in \partial B(x, \rho(x))$ by lemma 5, we have $\|a - x\| = \rho(x)$ and since $B(x, \rho(x)) \subset B(y, \rho(y))$, we know that $a \in \partial B(y, \rho(y))$ and so $\|a - y\| = \rho(y)$. Note that x is on the line segment $\overline{ay}$, for otherwise $B(x, \rho(x))$ would not be strictly contained in $B(y, \rho(y))$. Using this and the triangle inequality, $\|b - y\| \leq \|b - x\| + \|x - y\| = \|a - x\| + \|x - y\| = \|a - y\|$. Therefore $\|b - y\| \leq \|a - y\|$. If this were an equality, then a and b would be in $\partial B(y, \rho(y))$ and consequently we would not get $B(x, \rho(x))$ strictly contained in $B(y, \rho(y))$. Thus $\|b - y\| < \|a - y\| = \rho(y)$; i.e., $b \in B(y, \rho(y))$. Similarly, $b \notin B(y, \rho(y))$ implies $a \in B(y, \rho(y))$. So at least one of a and b is in $B(y, \rho(y))$, which is impossible since $B(y, \rho(y)) \cap A = \varnothing$, by lemma 4. $\square$

Proposition 3 *If* $\mathrm{FL}(a, b)$ *is nonempty and* $|\theta_a + \theta_b| \neq 180°$*, then* $\mathrm{FL}(a, b)$ *has nonzero length.*

PROOF: Suppose $\mathrm{FL}(a, b) \neq \varnothing$. Then there exists a point c in $\mathrm{FL}(a, b)$. Without loss of generality assume that both a and b lie on the ordinate a distance h from the origin, and that c lies on the nonnegative abscissa and so there exists $\tilde{c} \geq 0$ such that $c = (\tilde{c}, 0)$. Since both a and b are endpoints, denote the necessarily open curves which end there $\alpha : [0, 1] \to \mathbb{R}^2$ and $\beta : [0, 1] \to \mathbb{R}^2$ where $\alpha(0) = a$ and $\beta(0) = b$. For convenience let ϕ_a be the direction of the vector $\alpha'(0)$ and ϕ_b be the direction of $\beta'(0)$. Observe that $\phi_a = 90° - \theta_a$ and $\phi_b = -90° - \theta_b$ and so $|\theta_a + \theta_b| = 180°$ if and only if $|\phi_a + \phi_b| = 180°$. Consider the tangent at a of the maximal disc at c. Let ϕ_c be the angle, in the interval $[0, 90°)$, from the abscissa to this tangent. Since the maximal disc at c touches a and b, both $\alpha'(0)$ and $\beta'(0)$ must be directed away from that maximal disc, and therefore $\phi_a \in [\phi_c, \phi_c + 180°]$ and $\phi_b \in [-\phi_c - 180°, -\phi_c]$. Suppose $|\phi_a + \phi_b| \neq 180°$. If $\phi_a = \phi_c$, then $\phi_b \in (-\phi_c - 180°, -\phi_c]$, for if $\phi_b = -\phi_c - 180°$, we get $|\phi_a + \phi_b| = 180°$, which is impossible by our supposition. Similarly, if $\phi_a = \phi_c + 180°$, then $\phi_b \in [-\phi_c - 180°, -\phi_c)$. Thus we have two cases to consider: (i) $\phi_a \in [\phi_c, \phi_c + 180°)$,

$\phi_b \in (-\phi_c - 180°, -\phi_c]$; and (ii) $\phi_a \in (\phi_c, \phi_c + 180°]$, $\phi_b \in [-\phi_c - 180°, -\phi_c)$. This admittedly slight distinction is made to eliminate the pathological cases where either $\phi_a = \phi_c$ and $\phi_b = -\phi_c - 180°$ or $\phi_a = \phi_c + 180°$ and $\phi_b = -\phi_c$. In these cases, $\mathrm{FL}(a,b)$ is the singleton set $\{c\}$ and thus has zero length.

Let $F \triangleq (A \setminus (\alpha[0,1] \cup \beta[0,1])) \cup \alpha[s^*, 1] \cup \beta[s^*, 1]$ and note that $A = F \cup \alpha[0, s^*) \cup \beta[0, s^*)$, where s^* is defined using lemma 7. Consider a point q on the abscissa. Now,

$$\begin{aligned}\overline{B}(q, \|q-a\|) \cap A &= \overline{B}(q, \|q-a\|) \cap (F \cup \alpha[0, s^*) \cup \beta[0, s^*)) \\ &= (\overline{B}(q, \|q-a\|) \cap F) \cup (\overline{B}(q, \|q-a\|) \cap (\alpha[0, s^*) \cup \beta[0, s^*))).\end{aligned}$$

The first term in this union is empty for all q sufficiently close to c, by lemma 11. The second term is $\{a, b\}$ for some q arbitrarily close to c, by lemma 8. Thus $\overline{B}(q, \|q-a\|) \cap A = \{a, b\}$ for some $q \neq c$ on the abscissa and sufficiently close to c.

Using lemmas 12, 4, and 5,

$$\begin{aligned}\{a, b\} &= \overline{B}(q, \|q-a\|) \cap A = \overline{B}(q, \rho(q)) \cap A \\ &= (B(q, \rho(q)) \cap A) \cup (\partial B(q, \rho(q)) \cap A) = \varnothing \cup \pi(q).\end{aligned}$$

Combined with lemma 6, we know that q is a skeletal point with $\pi(q) = \{a, b\}$ and so $q \in \mathrm{FL}(a,b)$. Since $\mathrm{FL}(a,b)$ is connected, all the points on the abscissa between q and c are also in $\mathrm{FL}(a,b)$, and so the length of $\mathrm{FL}(a,b)$ is at least $\|q - c\| > 0$. $\square$

For the following two lemmas we must define angles ψ_a and ψ_b. Assume the conditions of case (i) (see the proof of proposition 3). Pick any $\epsilon \in (0, 90°)$. Let

$$\psi_a \triangleq \begin{cases} \max(\phi_c, \epsilon), & \text{if } \phi_a \in [\phi_c, \phi_c + 90°); \\ (\phi_a + \phi_c)/2, & \text{if } \phi_a \in [\phi_c + 90°, \phi_c + 180°). \end{cases} \tag{1}$$

Note that $\phi_c \leq \psi_a < \phi_c + 90°$ and $|\phi_a - \psi_a| < 90°$. Let

$$\psi_b \triangleq \begin{cases} \min(-\phi_c, -\epsilon), & \text{if } \phi_b \in (-\phi_c - 90°, -\phi_c]; \\ (\phi_b - \phi_c)/2, & \text{if } \phi_b \in (-\phi_c - 180°, -\phi_c - 90°]. \end{cases} \tag{2}$$

Observe similarly that $-\phi_c - 180^\circ < \psi_b \leq -\phi_c$ and $|\phi_b - \psi_b| < 90^\circ$. (For case (ii), we can make similar definitions of ψ_a and ψ_b.)

Lemma 7 *$\exists s^* > 0$ such that $\forall s \in [0, s^*), (\alpha(s) - a) \cdot (\cos\psi_a, \sin\psi_a) \geq 0$ and $(\beta(s) - b) \cdot (\cos\psi_b, \sin\psi_b) \geq 0$.*

PROOF: Because of our assumptions on A, we know that α is piecewise smooth and so we can apply Taylor's theorem on each of its coordinate functions α_1 and α_2 at 0:

$$\begin{aligned}
\alpha(s) &= (\alpha_1(s), \alpha_2(s)) \\
&= (\alpha_1(0) + \alpha_1'(0)s + O(s^2), \alpha_2(0) + \alpha_2'(0)s + O(s^2)) \\
&= a + \alpha'(0)s + (O(s^2), O(s^2)).
\end{aligned}$$

Thus,

$$(\alpha(s) - a) \cdot (\cos\psi_a, \sin\psi_a) = s\alpha'(0) \cdot (\cos\psi_a, \sin\psi_a) + O(s^2),$$

where the first term on the right hand side is strictly positive because $|\phi_a - \psi_a| < 90^\circ$. Therefore the sum is strictly positive for all $s \in (0, \tilde{s})$, for some small $\tilde{s} > 0$, and is exactly zero if $s = 0$.

Similarly, $\exists \hat{s} > 0$ such that $(\beta(s) - b) \cdot (\cos\psi_b, \sin\psi_b) \geq 0, \forall s \in [0, \hat{s})$. Let $s^* = \min(\tilde{s}, \hat{s})$ and the result follows. $\square$

Lemma 8 *$\exists q \neq c$, arbitrarily close to c and on the abscissa such that $\overline{B}(q, \|q - a\|) \cap (\alpha[0, s^*) \cup \beta[0, s^*)) = \{a, b\}$.*

PROOF: Assume the conditions of case (i) (see the proof of proposition 3) and define ψ_a and ψ_b using (1) and (2). Consider the closed half-plane H_a with a on its border and whose inward normal is $(\cos\psi_a, \sin\psi_a)$, or $H_a \triangleq \{(x, y) : y \geq \tan(\psi_a - 90^\circ)x + h\}$. Similarly, consider the closed half-plane H_b with b on its border and whose inward normal is $(\cos\psi_b, \sin\psi_b)$, or $H_b \triangleq \{(x, y) : y \leq \tan(\psi_b + 90^\circ)x - h\}$.

We first show that on the open left half-plane, $\overline{B}(q, \|q - a\|)$ is disjoint from both H_a and H_b for all q on the abscissa sufficiently close to c. Consider any point q on the abscissa, and so $(\tilde{q}, 0) \triangleq q$. For any point (x, y) in $\overline{B}(q, \|q - a\|)$, if $y^2 < (\tan(\psi_a - 90°)x + h)^2$, then $\{(x, y) : x < 0\} \cap \overline{B}(q, \|q - a\|) \cap H_a = \varnothing$. Since $\overline{B}(q, \|q - a\|) = \{(x, y) : (x - \tilde{q})^2 + y^2 \leq h^2 + \tilde{q}^2\}$,

$$
\begin{aligned}
& y^2 - (\tan(\psi_a - 90°)x + h)^2 \\
& \quad \leq \; h^2 + \tilde{q}^2 - (x - \tilde{q})^2 - (\tan(\psi_a - 90°)x + h)^2 \\
& \quad = \; -x^2(1 + \tan^2(\psi_a - 90°)) + 2x(\tilde{q} - h \tan(\psi_a - 90°)).
\end{aligned}
$$

The first term is strictly negative for $x < 0$. Since $\psi_a - 90° < \phi_c$ and $\tan \phi_c = \tilde{c}/h$, we know that $\tan(\psi_a - 90°) < \tilde{c}/h$ and so, for all $\tilde{q}$ sufficiently close to $\tilde{c}$, $\tan(\psi_a - 90°) < \tilde{q}/h$. Thus the second term will be negative for all $\tilde{q}$ sufficiently close to $\tilde{c}$. Therefore, for all q on the abscissa sufficiently close to c, $\{(x, y) : x < 0\} \cap \overline{B}(q, \|q - a\|) \cap H_a = \varnothing$. Similarly, for all q on the abscissa sufficiently close to c, $\{(x, y) : x < 0\} \cap \overline{B}(q, \|q - a\|) \cap H_b = \varnothing$. Combining, we conclude that, for all q on the abscissa sufficiently close to c, $\{(x, y) : x < 0\} \cap \overline{B}(q, \|q - a\|) \cap H_a \cap H_b = \varnothing$.

We now show that, for all q on the abscissa with $\tilde{q} < \tilde{c}$, $\{(x, y) : x \geq 0\} \cap \overline{B}(q, \|q - a\|) \subset \overline{B}(c, \|c - a\|)$. Fix a point $(x, y) \in \overline{B}(q, \|q - a\|)$, and so $y^2 \leq h^2 + \tilde{q}^2 - (x - \tilde{q})^2$. After a simple calculation we find that $(x - \tilde{c})^2 + y^2 - h^2 - \tilde{c}^2 \leq 2x(\tilde{q} - \tilde{c})$. If $x \geq 0$, then, for all $\tilde{q} < \tilde{c}$, $(x, y) \in \overline{B}(c, \|c - a\|)$, as required.

By lemma 7, we know that H_a contains $\alpha[0, s^*)$ and H_b contains $\beta[0, s^*)$. Putting these statements together, there exists q, arbitrarily close to c, on the abscissa, and to the left of c, such that:

$$
\begin{aligned}
& \overline{B}(q, \|q - a\|) \cap (\alpha[0, s^*) \cup \beta[0, s^*)) \\
& \quad = \; (\{(x, y) : x < 0\} \cap \overline{B}(q, \|q - a\|) \cap (\alpha[0, s^*) \cup \beta[0, s^*))) \\
& \qquad \cup (\{(x, y) : x \geq 0\} \cap \overline{B}(q, \|q - a\|) \cap (\alpha[0, s^*) \cup \beta[0, s^*))) \\
& \quad = \; \varnothing \cup \{a, b\},
\end{aligned}
$$

where we recall that $H_a \supset \alpha[0, s^*)$ and $H_b \supset \beta[0, s^*)$ and we use $\overline{B}(c, \|c - a\|) \cap (\alpha[0, s^*) \cup \beta[0, s^*)) = \{a, b\}$. The result follows for case (i). Case (ii) follows similarly. □

Lemma 9 *F is closed.*

PROOF: By the assumptions on A, α is a continuous mapping. By theorem 4.14 of Rudin [43], we know that $\alpha[s^*, 1]$ is compact and hence closed. Similarly, $\beta[s^*, 1]$ is closed. Since A is a *disjoint* union of traces (each of which is a closed set), and two of which are $\alpha[0, 1]$ and $\beta[0, 1]$, we know $A \setminus (\alpha[0, 1] \cup \beta[0, 1])$ is closed. Since finite unions of closed sets are closed we conclude F is closed. $\square$

Lemma 10 $F \cap \overline{B}(c, \|c - a\|) = \varnothing$.

PROOF: Since $c \in \mathrm{FL}(a, b)$, $\overline{B}(c, \rho(c)) \cap A = \overline{B}(c, \|c - a\|) \cap A = \{a, b\}$. But a is in $\alpha[0, s^*)$ and thus $a \notin F$, otherwise we would have intersecting or nonsimple curves whose traces were in A. Similarly, $b \notin F$, and the result follows. $\square$

Lemma 11 *For all q on the abscissa and sufficiently close to c, $\overline{B}(q, \|q - a\|) \cap F = \varnothing$.*

PROOF: By the triangle inequality, $\forall x \in \overline{B}(c, \|c - a\|), \forall y \in F, \forall q$, and $\forall z \in \overline{B}(q, \|q - a\|)$, we get $\|x - y\| \leq \|x - z\| + \|z - y\|$. Observe that $\overline{B}(c, \|c - a\|)$ is compact, F is closed (lemma 9), and $F \cap \overline{B}(c, \|c - a\|) = \varnothing$ (lemma 10). Hence there exists $\delta > 0$ such that, $\forall x \in \overline{B}(c, \|c - a\|)$ and $\forall y \in F$, $\|x - y\| > \delta$ [43, ex. 21 on p. 101]. Thus $\|z - y\| \geq \|x - y\| - \|x - z\| > \delta - \|x - z\|$. Now for all q sufficiently close to c, there exists $x \in \overline{B}(c, \|c - a\|)$ such that $\|x - z\| < \delta/2$, say. Thus for all q on the abscissa and sufficiently close to c, for all z in $\overline{B}(q, \|q - a\|)$, and for all y in F, $\|z - y\| > \delta/2$. Therefore $\overline{B}(q, \|q - a\|)$ and F are disjoint. $\square$

Lemma 12 *If $\overline{B}(q, \|q - a\|) \cap A = \{a, b\}$, then $\rho(q) = \|q - a\|$.*

PROOF: Since $a \in A, \rho(q) \leq \|q - a\|$. If $\rho(q) < \|q - a\|$, then $\exists p \in A$ such that $\|q - a\| > \|q - p\| \geq \rho(q)$, and therefore $p \in \overline{B}(q, \|q - a\|)$. However since $\overline{B}(q, \|q - a\|) \cap A = \{a, b\}$, p must be either a or b, which is a contradiction because $\|q - a\| = \|q - b\|$. $\square$

Proposition 4 *The midpoint between a and b lies in* $\mathrm{FL}(a, b)$ *if and only if r has a local minimum in* $\mathrm{FL}(a, b)$.

PROOF: Let L denote the perpendicular bisector of the line segment $\bar{ab}$. Clearly the midpoint $m \triangleq \frac{a+b}{2}$ of $\bar{ab}$ is in L. For $q \in L$, $||a - q||$ is minimized when the vector $a - q$ is perpendicular to L, but that occurs exactly when $q = m$. If $m \in S$, then $r(q) = ||a - q||$ takes on a local minimum at $q = m \in \mathrm{FL}(a, b)$. $\square$

Acknowledgements: Many thanks to Yves Bérube-Lausière, Shamez Alibhai, Ohad Ben-Shahar, Mike Kelly. In addition, we appreciate the suggestions of the reviewers, in particular for Proposition 4.

References

[1] N. Ahuja and M. Tuceryan. Extraction of early perceptual structure in dot patterns: integrating region, boundary, and component gestalt. *Computer Vision, Graphics, and Image Processing*, 48:304–356, 1989.

[2] A. Amir and M. Lindenbaum. A generic grouping algorithm and its quantitative analysis. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 20:168–185, 1998.

[3] H. Blum. Biological shape and visual science. *J. Theor. Biol.*, 38:205–287, 1973.

[4] M. Boldt, R. Weiss, and E. M. Riseman. Token-based extraction of straight lines. *IEEE Trans. on Systems, Man, and Cybernetics*, 19:1581–1594, 1989.

[5] A. L. Bregman. Asking the "what for" question in auditory perception. In M. Kubovy and J. R. Pomerantz, editors, *Perceptual Organization.* Erlbaum, Hillsdale, NJ, 1981.

[6] R. Brockett and P. Maragos. Evolution equations for continuous-scale morphology. In *Proceedings of the IEEE Conference on Acoustics, Speech and Signal Processing*, San Francisco, CA, March 1992.

[7] L. Calabi and W. E. Hartnett. Shape recognition, prairie fires, convex deficiencies and skeletons. *American Mathematical Monthly*, 75:335–342, 1968.

[8] S. Casadei and S. Mitter. Hierarchical image segmentation - part i: Detection of regular curves in a vector graph. *International Journal of Computer Vision*, 27:71–100, 1998.

[9] I. J. Cox, J. M. Rehg, and S. Hingorani. A Bayesian Multiple-Hypothesis Approach to Edge Grouping and Contour Segmentation. *International Journal of Computer Vision*, 11:5–24, 1993.

[10] C. David and S. W. Zucker. Potentials, Valleys, and Dynamic Global Coverings. *International Journal of Computer Vision*, 5:219–238, 1990.

[11] B. Dubuc and S. W. Zucker. Indexing visual representations through the complexity map. In *Fifth International Conference on Computer Vision*, pages 142–149. IEEE, 1995.

[12] J. Elder and S. W. Zucker. The effect of contour closure on the rapid discrimination of two-dimensional shapes. *Vision Research*, 33(7):981–991, 1993.

[13] J. H. Elder. *The Visual Computation of Bounding Contours.* PhD thesis, McGill University, 1995.

[14] J. H. Elder and S. W. Zucker. Computing Contour Closure. In *Proceedings, 4th European Conference on Computer Vision*, pages 399–412, 1996.

[15] D. Geiger, K. Kumaran, and L. Parida. Visual organization for figure/ground separation. In *Proc. IEEE Conf. on Computer Vision and Pattern Recognition*, pages 155–160, 1996.

[16] G. Guy and G. Medioni. Inferring global perceptual contours from local features. *IJCV*, 20(1/2):113–133, 1996.

[17] J. Hadamard. *Lectures on the Cauchy problem in linear partial differential equations.* Yale U. P., New Haven, 1923.

[18] D. D. Hoffman and W. A. Richards. Parts of recognition. *Cognition*, 18:65–96, 1985.

[19] B. K. P. Horn. The Curve of Least Energy. *ACM Transactions on Mathematical Software*, 9:441–460, 1983.

[20] M. Ilg. Knowledge-based understanding of road maps and other line images. In *Proceedings, 10th International Conference on Pattern Recognition*, volume 1, pages 282–284, 1990.

[21] L. A. Iverson and S. W. Zucker. Logical/linear operators for image curves. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 17(10):982–996, 1995.

[22] D. Jacobs. Grouping for recognition. MIT AI Lab Memo No. 1177, MIT, 1989.

[23] G. Kanizsa. *Organization in Vision.* Praeger, New York, 1979.

[24] P. J. Kellman and T. F. Shipley. A theory of visual interpolation in object perception. *Cognitive Psychology*, 23:141–221, 1991.

[25] B. B. Kimia, A. R. Tannenbaum, and S. W. Zucker. Shapes, shocks, and deformations, I: The components of shape and the reaction-diffusion space. *International Journal of Computer Vision*, 15:189–224, 1995.

[26] J. J. Koenderink. *Solid Shape.* MIT Press, 1990.

[27] K. Koffka. *Principles of Gestalt Psychology.* Harcourt, Brace & World, New York, 1963.

[28] P. D. Lax. *Hyperbolic Systems of Conservation Laws and the Mathematical Theory of Shock Waves.* SIAM Regional Conference series in Applied Mathematics, Philadelphia, 1973.

[29] A. B. Lee and D. Mumford. An occlusion model generating scale invariant images. In *Workshop on Statistical and Computational Theories of Vision*, Fort Collins, Colorado, 1999.

[30] T. Leung and J. Malik. Contour continuity in region based image segmentation. In *Proceedings, 5th European Conference on Computer Vision*, 1998.

[31] G. Lohmann. A new method of extracting closed contours using maximal discs. In *Proc. of 6th Intl. Conf. on Computer Analysis of Images and Patterns*, volume 970 of *Lecture Notes in Computer Science*, pages 472–479, Berlin, 1995. Springer-Verlag.

[32] D. Lowe. *Perceptual Organization and Visual Recognition.* Kluwer Academic Press, The Netherlands, 1985.

[33] R. Mohan and R. Nevatia. Perceptual Organization for Scene Segmentation and Description. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 14(6):616–635, 1992.

[34] D. Mumford. *Algebraic Geometry and Its Applications*, chapter Elastica and Computer Vision, pages 491–506. Springer-Verlag, 1994.

[35] K. Nakayama and S. Shimojo. Experiencing and perceiving visual surfaces. *Science*, 257:1357–1363, 1992.

[36] M. Nitzberg, D. Mumford, and T. Shiota. *Filtering, Segmentation and Depth.* Springer-Verlag, Berlin, 1993.

[37] A. Okabe, B. Boots, and K. Sugihara. *Spatial Tessellations.* Probability and Mathematical Statistics. Wiley, Chichester, NY, 1992.

[38] S. Osher and J. Sethian. Fronts propagating with curvature dependent speed: Algorithms based on Hamilton-Jacobi formulations. *Journal of Computational Physics*, 79:12–49, 1988.

[39] 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.

[40] S. Petry and G. E. Meyer, editors. *The Perception of Illusory Contours.* Springer-Verlag, New York, 1987.

[41] R. A. Rensink and J. T. Enns. Early completion of occluded objects. *Vision Research*, 38:2489–2505, 1998.

[42] P. L. Rosin. Grouping curved lines. In *Proceedings, British Machine Vision Conference*, 1994.

[43] W. Rudin. *Principles of Mathematical Analysis.* McGraw-Hill, New York, third edition, 1976.

[44] S. Sarkar and K. L. Boyer. Quantitative measures of change based on feature organization: Eigenvalues and eigenvectors. *Computer Vision and Image Understanding*, 71(1):110–136, 1998.

[45] A. Sha'ashua and S. Ullman. Structural saliency: The detection of globally salient structures using a locally connected network. In *Proceedings, Second International Conference on Computer Vision.* IEEE Computer Society Press, 1988.

[46] D. Shaked and A. M. Bruckstein. On symmetry axes and boundary curves. In C. Arcelli, editor, *Aspects of Visual Form Processing*, pages 497–506, London, 1994. World Scientific.

[47] R. Shapley and J. Gordon. The existence of interpolated illusory contours depends on contrast and spatial separation. In S. Petry and G. E. Meyer, editors, *The Perception of Illusory Contours*, chapter 11, pages 109–115. Springer-Verlag, 1987.

[48] J. Shi and J. Malik. Normalized cuts and image segmentation. In *Proceedings, Computer Vison and Pattern Recognition*, 1997.

[49] K. Siddiqi and B. B. Kimia. A shock grammar for recognition. Technical Report LEMS 143, LEMS, Brown University, September 1995.

[50] Z. S. G. Tari, J. Shah, and H. Pien. Extraction of shape skeletons from grayscale images. *Computer Vision and Image Understanding*, 66(2):133–146, 1997.

[51] H. Tek, P. A. Stoll, and B. B. Kimia. Shocks from images: propagation of orientation elements. In *CVPR*, pages 839–845, 1997.

[52] S. Ullman. Filling-in gaps: The shape of subjective contours and a model for their generation. *Biological Cybernetics*, 25:1–6, 1976.

[53] L. Williams and A. Hanson. Perceptual completion of occluded surfaces. *Computer Vision and Image Understanding*, 64(1):1–20, 1996.

[54] L. Williams and D. Jacobs. Local parallel computation of stochastic completion fields. *Neural Computation*, 9(4):859–881, 1997.

[55] L. Williams and D. Jacobs. Stochastic completion fields: A neural model of illusory contour shape and salience. *Neural Computation*, 9(4):837–858, 1997.

[56] L. Williams, T. Wang, and K. Thornber. Computing stochastic completion fields in linear-time using a resolution pyramid. In *Proc. of 7th Intl. Conf. on Computer Analysis of Images and Patterns*, Kiel, Germany, 1997.

[57] L. R. Williams. *Perceptual Completion of Occluded Surfaces.* PhD thesis, University of Massechusettes at Amherst, 1994.

[58] A. Yl a-J a aski and F. Ade. Grouping symmetrical structures for object segmentation and description. *Computer Vision and Image Understanding*, 63(3):399–417, 1996.

[59] A. Zisserman, J. Mundy, D. Forsyth, and J. Liu. Class-based grouping in perspective images. In *Proceedings, Fifth International Conference in Computer Vision*, pages 183–187. IEEE, 1995.

[60] S. W. Zucker. The diversity of perceptual grouping. In M. Arbib and A. Hanson, editors, *Vision, Brain, and Cooperative Computation.* MIT Press, Cambridge, 1986.