Journal of Experimental Psychology: General

Constant Curvature Segments as Building Blocks of 2D Shape Representation

Nicholas Baker, Patrick Garrigan, and Philip J. Kellman

Online First Publication, December 17, 2020. http://dx.doi.org/10.1037/xge0001007

AMERICAN PSYCHOLOGICAL ASSOCIATION

Journal of Experimental Psychology: General

ISSN: 0096-3445

https://doi.org/10.1037/xge0001007

Constant Curvature Segments as Building Blocks of 2D Shape Representation

Nicholas Baker[1], Patrick Garrigan[2], and Philip J. Kellman[1]
[1] Department of Psychology, University of California, Los Angeles
[2] Department of Psychology, Saint Joseph's University

How the visual system represents shape, and how shape representations might be computed by neural mechanisms, are fundamental and unanswered questions. Here, we investigated the hypothesis that 2-dimensional (2D) contour shapes are encoded structurally, as sets of connected constant curvature segments. We report 3 experiments investigating constant curvature segments as fundamental units of contour shape representations in human perception. Our results showed better performance in a path detection paradigm for constant curvature targets, as compared with locally matched targets that lacked this global regularity (Experiment 1), and that participants can learn to segment contours into constant curvature parts with different curvature values, but not into similarly different parts with linearly increasing curvatures (Experiment 2). We propose a neurally plausible model of contour shape representation based on constant curvature, built from oriented units known to exist in early cortical areas, and we confirmed the model's prediction that changes to the angular extent of a segment will be easier to detect than changes to relative curvature (Experiment 3). Together, these findings suggest the human visual system is specially adapted to detect and encode regions of constant curvature and support the notion that constant curvature segments are the building blocks from which abstract contour shape representations are composed.

Keywords: visual perception, curvature, contour shape, shape perception, shape representation

Among the most important problems of visual perception is how we perceive and encode the shapes of objects. Shape representations capture critical affordances of objects, allowing us to interact with them through reaching and grasping (Lederman & Wing, 2003), infer their functional properties (Graham et al., 2004; Welder & Graham, 2001), and recognize them, both within (Collin et al., 2004) and across basic categories (Murphy & Brownell, 1985; Rossion & Pourtois, 2004). Despite their importance, surprisingly little is known about the kind of information contained in shape representations, nor about the neural mechanisms by which they are encoded.

Shape is inherently a relational notion. Whether conveyed by the positions of the edges of objects or by discrete perceptible elements, shape inheres in the relations of the positions of elements, not in the positions or elements themselves. Having a certain shape does not require being in a particular place or having any particular elements as constituents, nor does it involve properties of the elements other than their spatial (and sometimes temporal) relations. The Gestalt psychologists were most eloquent in pointing out these ideas about shape and also in emphasizing the centrality of shape and relations in perception (Koffka, 1935; Kohler, 1929; Wertheimer, 1923).

Experimentally, the importance of shape and configurational relations has been shown in many ways. In perceptual processing, for example, properties of configurations are often more rapidly accessed than properties of individual elements (Pomerantz & Portillo, 2011; Pomerantz et al., 1977), such that judgments relating to parts of the display are faster and easier when these parts are embedded in configurational contexts. Shape representations are conferred even on sparse, separated, local elements, and such abstract shapes are retained in perceptual representations even when the constituent elements are not. For example, Baker and Kellman (2018), using shapes made from separated elements, found that observers had no sensitivity to changes in the spatial positions of elements so long as their overall configuration was preserved.

Implicit in these and other aspects of the phenomenology and perceptual processing of shapes is the idea that shape representations are *abstract*. Abstraction is a complicated notion, having several related meanings (see Barsalou & Wiemer-Hastings, 2005, for useful discussion). In our use of "abstract" in the realm of shape perception, we intend three related ideas. The first is encompassed by the idea above that shape is a relational notion. That is, by abstract, we mean at least that relevant information consists in relationships defined over, but not by, lower order constituents; such relationships can be described as binding the value of a variable (Marcus, 2001; Overlan et al., 2017). For instance, to be a square does not mean that a side of the form has to be of a certain length, but that the length of one side must equal the length of any other side, formally expressible as, for any two sides a and b,

Nicholas Baker https://orcid.org/0000-0002-0673-2486
Nicholas Baker is now at Centre for Vision Research, York University.
Portions of this work have been presented at annual meetings of the Vision Science Society and the Configural Processing Consortium.
Correspondence concerning this article should be addressed to Philip J. Kellman, Department of Psychology, University of California, Los Angeles, 405 Hilgard Avenue, Los Angeles, CA 90095-1563, United States. Email: Kellman@cognet.ucla.edu

length(a) is equal to length(b). Intuitively, we recognize shape representations as abstract when we notice that a cloud appears to resemble a dog, or when we notice that two objects of different sizes, composition, and orientation share the same shape. Our perceptual representations of shape allow these "matching" experiences despite radically different contexts or constituent elements. A second, related, idea is that at least some shape descriptions, including, arguably, those used in the brain, capture information economically (Attneave, 1954); they comprise a summary description from which much specific stimulus information has been discarded. A third aspect of abstraction in human shape representation is more or less the converse: The representation is abstract in adding something that was not present in the stimulus. Baker and Kellman (2018) studied abstract shape perception experimentally using separated dot elements around virtual contours of unfamiliar, smooth, closed two-dimensional (2D) shapes. They found that, beyond a very brief interval after stimulus offset, encoding of specific elements was poor or nonexistent. In contrast, shape representations were encoded that supported accurate same/different judgments, across displays, despite transformations of position, orientation, and scale. Such results imply a representation that has captured relations among the inputs while discarding the concrete values of the inputs. Moreover, in these studies, no continuous contour shape information was actually given in the stimulus. The dots used in each display could have been connected in a virtually unlimited number of ways (or not at all). The particular shape representations that supported task performance were abstract in (a) being derived from relations; (b) being more economical descriptions in that the input elements were not stored; and (c) in supplying connections across gaps in the physically specified input.

These criteria imply that abstract shape representations are *symbolic* representations, in that they designate properties of material objects in the world (cf., Simon & Newell, 1976). One of the deepest complexities of visual perception is that it involves a transition to symbolic descriptions of the environment from initially *subsymbolic* inputs. Early processing in the visual pathway involves units that respond to light or contrast. Encoding these properties of light energy, while crucial to vision, is not the *goal* of vision; rather, representing material properties of the world, such as objects, arrangements, and events are goals of vision (Gibson, 1966, 1979; Marr, 1982), and extraction of contrast occurs in the service of the development of richer descriptions (Neri, 2018). The transition in visual perception from subsymbolic to symbolic coding largely corresponds to the distinction between responses to properties of incident light *energy*, as in the activation of a retinal photoreceptor or an oriented contrast detector in V1, versus representing properties of *matter*—objects in the world. Activations of the former type change markedly with fluctuations in illumination, small changes in observer position, orientation, distance, and so forth. They are also transient, in that they ordinarily appear not to be preserved in more durable encodings (e.g., Baker & Kellman, 2018; Sperling, 1960), serving primarily to facilitate the extraction of subsequent representations. On the other hand, the perceptual description that one is seeing a rectangular table is symbolic. This sort of representation is more durable and relatively invariant across changes in illumination, observer position, and so on. Representations of properties of objects, such as their unity, continuity, shape, and material composition, are symbolic representations. The transition from subsymbolic to symbolic processing remains deeply mysterious in visual perception (Kellman et al., 2013), and most research occurs on one side or the other of this divide. Although the present work focuses specifically on understanding the representation of contour shape, it also has a larger purpose of using this domain as an example and existence proof of how the visual system may obtain symbolic descriptions from initially subsymbolic encodings.

Structural Descriptions in Vision

Proposals for representational schemes to describe shape have varied across different aspects of shape perception and recognition in vision, but they have commonalities. They tend to be *object-centered*, that is, parts of an object are represented with respect to an origin and axes centered on the object (Marr, 1982). Object-centered representations that might afford the flexibility and generality of human shape perception likely involve *structural descriptions*. Structural descriptions represent object shape in terms of a limited, predefined set of parts, called primitives, and the spatial relations among them. Whether structural descriptions exclusively account for human performance as assessed in psychophysical experiments has been a subject of debate. Both structural descriptions and view-based approaches may play a role in human perception. In particular, viewpoint-based models of object recognition provide a parsimonious explanation for a variety of viewpoint-dependent empirical results (Tarr & Bulthoff, 1995), and have achieved considerable success in terms of implementation and application to real images (e.g., Ullman, 2007).

Structural descriptions do, however, appear necessary to account for some human object recognition capabilities. Similar considerations apply to artificial systems. Recent image-based recognition systems based on deep convolutional neural networks show impressive performance on object categorization tasks, but fail catastrophically in situations where overall shape information is important for object classification (Baker et al., 2018). Such shortcomings of image-based approaches appear to derive from the lack of explicitly encoded structure.

In structural description schemes, basic elements (primitives) provide a means for representing a large class of stimuli (Marr & Nishihara, 1978). With appropriate primitives, a structural description scheme can encode a wide variety of shapes and capture ecologically relevant similarities and differences among objects or patterns. Particular sorts of primitives may be valuable because they can be used to represent frequently encountered stimulus patterns, because they can be combined to represent a large class of stimuli, or because they are easily extracted from the visual environment. Various representational schemes with these properties have been used to model three-dimensional (3D) shape (Marr & Nishihara, 1978), 3D shape as captured by 2D, nonaccidental properties (Biederman, 1987), and 2D shape (Feldman & Singh, 2006).

Although structural descriptions have most notably been applied to problems of 3D or 2D shapes of objects, the issues that call for structural representations are also challenges for more basic shape representation. In this work, we focus on perhaps the most elementary shape problem: perceiving 2D contour shape. Like 3D volumetric shapes, 2D contours can be encoded, stored, and later compared with other 2D contour shapes. Humans can recognize

2D contours as having the same or different shapes across transformations of translation, rotation, and scaling (Baker & Kellman, 2018). The shape of one 2D contour can be compared with the shape of part of another 2D contour. All of these abilities suggest that human visual perception builds structural descriptions of 2D contour shapes. Figure 1 illustrates some of these 2D contour perception abilities. Before looking at the caption, note the difference in shape between (A) and (B); see which of these shapes matches the scaled, rotated shape in (C); and, for the rotated fragment in (D), identify from which shape, (A) or (B), it has been extracted. These capacities would be difficult to explain without some compositional description of the shape.

Structural description models have been proposed to encode contour shapes. Modern symmetry-based models (e.g., Feldman & Singh, 2006; Rezanejad & Siddiqi, 2013) represent 2D shape as a set of axial branches, including structural primitives to capture important contour information and abstracting over inessential contour features. One problem with symmetry-based models is that they are not well suited to handle open contours. To the extent that algorithms can extract a skeletal description from an open contour at all, the skeleton that results from a contour fragment would bear no similarity to any part of the skeletal description obtained from the closed contour. Skeletal representations, then, could not explain performance on a partial shape matching task using stimuli such as those in Figure 1B and Figure 1D. Behavioral evidence for symmetry-based models has also been limited, although some work has found evidence for better visual acuity along a shape's principal axis (Kovacs & Julesz, 1993) and that skeletal representations accurately predict people's judgments of two shapes' similarity (Ayzenberg & Lourenco, 2019; Lowet et al., 2018).

Contour-based structural shape models have been proposed. Some, such as Fourier descriptors (e.g., Zahn & Roskies, 1972; Zhao & Belkasim, 2012) and active contour models (Kass et al., 1988) have components that are not localized in space and are therefore not robust to partial occlusion. Though compositional, these models are arguably not structural, as each component applies to the whole contour. Other contour-based shape theories that satisfy this locality constraint involve deformation of an embryonic shape by the addition of morphing primitives (e.g., Dubinskiy & Zhu, 2003; Elder et al., 2013). These impressive models, originating from computer vision, use sophisticated mathematical tools to capture shape representation. As engineering solutions, these systems might work well, but no explanation is given for how they connect to outputs of subsymbolic systems. The constant curvature theory we propose below has a more straightforward connection to outputs from early visual areas.

Figure 1
Examples of Structure in Contour Shape Perception

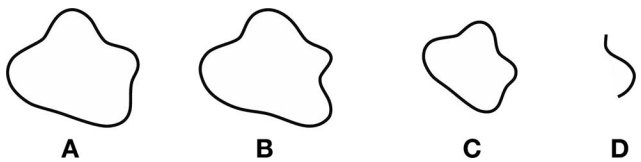

Note. Despite overall similarity of size and shape, shapes (A) and (B) can be readily distinguished. Despite rotation and scaling, it is apparent that (C) shares the same shape as (A). Although its orientation has been changed, the fragment in (D) can be seen to match part of shape (B).

Other lines of investigation relevant to shape perception and representation have focused on how perceivers segment contours or planar shapes into parts (Wertheimer, 1923; see De Winter & Wagemans, 2006 for a review) or the relevance of concavities and convexities in shape processing tasks (e.g., Barenholtz et al., 2003; Bertamini, 2001; Schmidtmann et al., 2015). These efforts have implicated various shape features as relevant, with wide variation across tasks. In general, the approach to contour shape representation presented here is complementary to these efforts. Research on part segmentations of objects or contours takes as a given the veridical representation of the contours themselves. It does not address how representations of contour shape form initially, nor does it suggest ways in which a represented shape may differ from the stimulus presented (Garrigan & Kellman, 2011). Consistent with this difference in focus, approaches to part segmentation also do not build from early visual detectors to obtain a representation. The issues we attempt to address here, then, are basic to visual coding and other visual tasks, yet are relatively unexplored. It is likely that contour shape representation as investigated here is a more basic rendering from the stimulus input to a representational scheme than are the more ecologically relevant segmentations of objects into parts. Both get at important aspects of perception and cognition of shape, but the former is more or less presupposed by the latter. We return to these issues in the General Discussion.

Constant Curvature Shape Representation

A plausible model of contour shape representation in human visual perception should satisfy at least three criteria. First, it should efficiently encode ecologically useful information about common objects' physical boundaries, as object perception, recognition, and representation are perhaps the most important functions of shape descriptions. Although 3D shape representations are surely important, we are able to represent shape from single views or pictures. Consequently, contours should be represented with sufficient precision to distinguish individuals, such as a silhouette of my dog from yours, or my coffee mug from yours, but also structurally descriptive in ways that capture similarities, such as the similarities among all dogs and among all mugs. This structural description will include some basic set of primitives that can be combined to represent an unlimited set of actual contour shapes. Moreover, the representation should be abstract enough to capture shape invariance across transformations of position, orientation, and scale (Baker & Kellman, 2018). The generative nature of structural descriptions relates to the second criterion: The representations must be compact. The set of activated neural units in V1 produced by presentation of an object must in some sense contain the information from which other representations can be derived, but as representations become more abstract, they should become more efficient and compact (Farah et al., 1994). Data compression in the encoding of a contour's shape is essential to efficient processing of visual information and to support some degree of shape invariance despite small, local variations along its contour (Barenholtz et al., 2003; Bell et al., 2007). Third, models of shape description should be consistent with empirical evidence about human shape processing and with what is known about neural coding. Models that are substantively different make different behavioral predictions about how humans perceive, interact, and make judgments about shapes, and those models whose predictions

better align with human behavior should be given greater consideration. Likewise, models that connect to outputs of known neural mechanisms may have greater plausibility.

One way the visual system might encode structural descriptions of 2D contour shape that fit the criteria above is by partitioning contiguous contours into regions of similar curvature, encoding these regions as segments of constant curvature, and representing the full shape in terms of these constant curvature parts and their spatial relations. This would entail finding regions along the contour with relatively little curvature variance and segmenting them from adjacent regions in which the distribution of curvatures differs. These regions might then be recoded into segments with a single curvature value, abstracting over the variation within the contour region. A symbolic code would result, in which the set of contour points within each segment would be described as a single curvature value, position, orientation, and angular extent (see Figure 2). The full, abstract representation of an object's contour shape would then be composed from the piecewise connections between constant curvature segments.

Under this hypothesis, segments of constant curvature (CC) are the elementary primitives from which contour shape representations are constructed. One advantage of the CC hypothesis is that there is a plausible way for such symbolic representations to be derived from early, subsymbolic responses in the visual system (Kellman et al., 2013). Classic research in early cortical visual areas has found receptive fields that are sensitive to oriented luminance contrast in local regions. These neural units occur across retinal locations, spanning various orientations and a range of spatial frequencies (De Valois & De Valois, 1980; Hubel & Wiesel, 1968). The sensitivity profiles of these units are well described by Gabor functions, each of which is a product of a periodic function of luminance in some orientation and a Gaussian window (Kulikowski & Bishop, 1981; Marĉelja, 1980). An important property of these detectors is they are isotropic across scales (Ringach, 2002). Detectors of various sizes have similar shapes, and receptive fields contain approximately the same number of cycles of the relevant spatial frequency across sizes. This occurs because receptive field size varies inversely with spatial frequency; receptive fields for high spatial frequencies are smaller than for low spatial frequencies. The filtering of images with local oriented detectors across multiple scales captures information that can be used to estimate curvature and also turns out to enable extraction of constant curvature representations of the same shape across differing stimulus sizes.

Figure 2

An Example of an Arclet Unit Built From Relations of Cortical Oriented Units

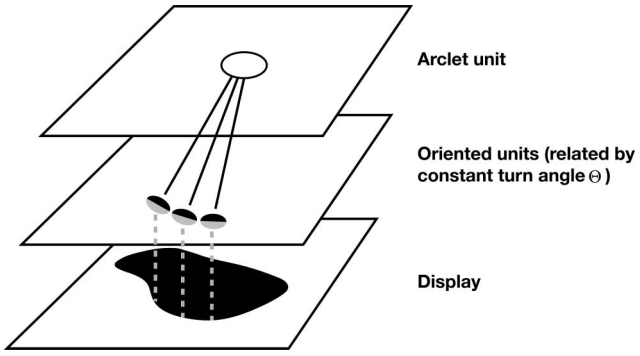

Note. An arclet is a higher-order detector that is activated when two or more oriented units forming a collinear or cocircular path are simultaneously activated. At the bottom of the figure is the viewed object. The object activates sets of oriented units (shown as Gabor filters) in early cortical areas. (Dashed lines indicate relations of the 2D shape contour to the spatially best-fitting filters.) Adjacent arclets having the same (or similar) turning angle and scale that respond along a contour segment are linked to each other so that longer segments of constant curvature can be detected. Arclets are hypothesized to exist across a range of turning angles and scales, and a given contour segment may activate arclets at different scales (see text).

Theoretical Model of Constant Curvature Shape Representation

In earlier research (Garrigan & Kellman, 2011; Kellman & Garrigan, 2007; Kellman et al., 2013) we proposed a general scheme for how this might work. We hypothesized the existence of higher-order neural circuits, which we call *arclets*, that respond to cocircular or nearly cocircular oriented units, that is, adjacent oriented units that are linked by constant *turning angles*. Turning angle (the difference in orientation between contrast detectors in an arclet), though mathematically different from a continuous estimate of curvature, serves as an approximation of curvature between lower-order detectors.

An arclet is essentially an "and" gate that is connected to a pair (or possibly more) of adjacent oriented units that have a cocircular relation. It is activated when two or more of these oriented units are concurrently activated. Arclets vary in their preferred curvature via the turning angle relating their input orientation sensitive units. They also vary in spatial frequency. As each arclet is comprised of two input units of the same spatial frequency, the set of arclets will span not just a range of turning angles (curvatures) but will also span the range of spatial frequencies present in cortical oriented units. We envision that, like the oriented units from which they are comprised, arclets exist and operate in parallel at positions and orientations across much or all of the visual field.

In past work, we have proposed a computational model of how this encoding could work, and we refer the reader to that work for details (Garrigan & Kellman, 2011). Current efforts involve implementing a model of curvature encoding as a fully specified neural model (Baker, 2020). As shown in Figure 3, an arclet is activated if a pair of oriented units forming a collinear or cocircular path are simultaneously activated. At the bottom of the figure is the viewed object. The object activates sets of oriented units (shown as Gabor filters) in early cortical areas. Adjacent arclets having the same (or similar) turning angle and scale that respond along a contour segment are linked to each other so that longer segments of constant curvature can be detected. A given contour segment may activate arclets at different scales (i.e., the receptive field size of the oriented contrast detectors to which the arclet is connected). For example, the border of a large object may be well fit by arclets that are built from large (low spatial frequency) detectors, but any perceptible contour will also activate arclets at the smallest scale, as well as intermediate scales. For a given segment of approximately constant curvature, arclets of different scales and turning angles may fit to differing degrees. It is likely

Figure 3

Scale Invariant and Scale Specific Shape Coding From Arclets

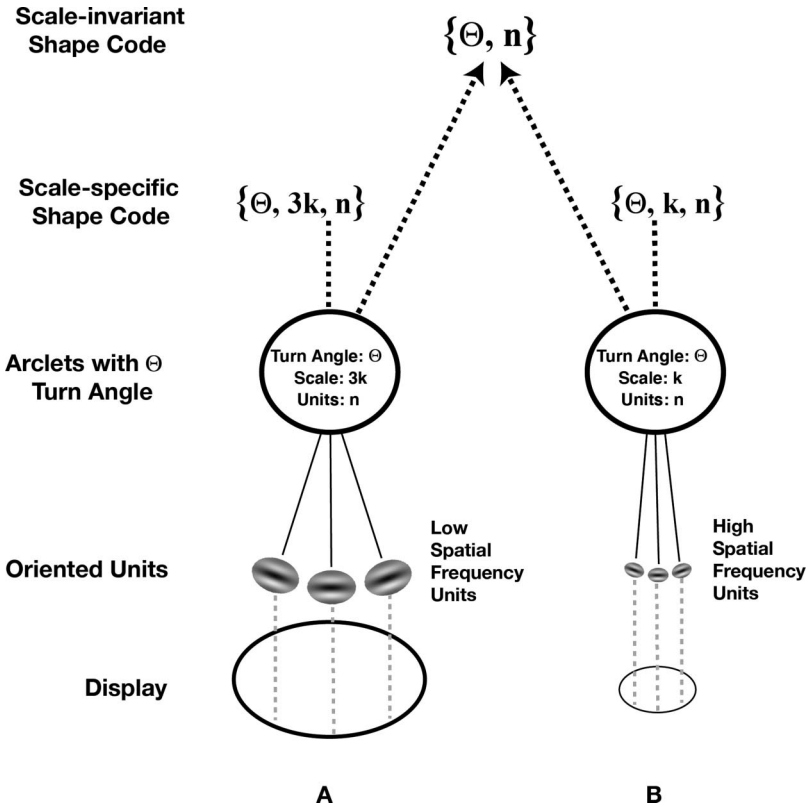

Note. (A) Large ellipse display. A contour fragment of approximately constant curvature initially activates oriented units along its boundary. Here, the largest adequately fitting arclets for the segment are given as scale 3k. The turning angle along this segment is Θ, and the extent of the constant curvature segment best responded to by arclets of this turning angle and scale is *n* units. (B) Small ellipse display. Arclet units are activated by small coaxial oriented units along the corresponding contour segment as in (A) that are related by a constant turning angle Θ. Here, the largest adequately fitting arclet has scale k, turning angle Θ, and extent of *n* units. A scale-specific representation of this segment of the ellipse is given by the three parameters of the fitting arclets: turning angle, scale, and number of oriented units comprising the extent of the segment. This scale-specific representation will differ for the large and small ellipse. Encoding at the largest adequately fitting scale makes available a scale-invariant representation. Omitting the scale parameter, the two segments of the two ellipses have the same shape because they match on the two parameters of turning angle and number of segments (of the largest adequately fitting arclets in each case). The visual system's use of isotropic operators at different scales and the detection of curvature from sets of straight oriented units related by constant turning angles makes this scale invariant code available without special computation (e.g., normalization). Other parts of the contour will have other best-fitting arclets to signal approximately constant curvature segments. The complete contour representation of each ellipse consists of several joined segments of approximately constant curvature (see text.)

that information from multiple scales is used in contour representations, but an important property of this scheme is that it identifies the *largest scale* that produces an adequately good fit (below some threshold of error) to the actual contour segment. These largest arclets that fit the curve contribute three pieces of information to a symbolic code of contour shape: the scale (spatial frequency) of the oriented units, the turning angle relating them (Θ deg in the example given), and the number of oriented units (encoding the length of the constant curvature segment).

The encoding of a constant curvature segment extends along a contour until a transition zone, where the fit for that curvature falls below some accuracy threshold, so that a different curvature arclet provides a better fit. The complete shape representation consists of a set of arclets describing adjacent regions along the contour. (For a working computational model of this scheme, see Garrigan & Kellman, 2011; for more details of the neural model, see Baker, 2020).

A remarkable property emerges from the fact that arclets operate at multiple spatial scales and are derived from sets of isotropic

oriented units at multiple scales. In such a scheme, scale invariant shape is directly available. Consider two ellipses, one with a major axis three times as long as the other. How do we explain the perceptually obvious fact that they have the same shape? (Shape for human observers appears to be a scale-invariant notion.) The actual values of mathematical curvature for these two objects (the change in contour orientation per unit arc length) differ for the large and small ellipses at every corresponding point. So, what accounts for our seeing them as having the same shape? The most common methods used in biological and computational vision models for equating shape across entities of differing size is to take some global measurement (e.g., the longest axis of object) and use it as a normalizing factor. For the ellipses, we find that if we divide the length of the major radius of the larger by the length of the major radius of the smaller, that scale factor obtained will also be the scale factor relating the curvatures at any point. Up to this scale factor, the curvatures at every point will be the same.

In a representational scheme based on arclets (or equivalent), normalization of this sort is not required. Because the visual system employs multiscale filtering and encodes curves by relations of oriented units, such normalization is essentially built in. The key properties are these: (a) As long as all elements within each arclet are of equal size, all arclets based on the same turning angle between oriented elements represent the same scale-invariant shape, that is, shape pieces that differ only by a scalar. (b) For a curved segment of any size, activating the arclet of largest scale that fits the curve signals yields two values: a turning angle (curvature), and the number of arclet units needed to encode the length of that constant curvature segment (*angular extent*). Any smaller or larger version of that segment will have a different scale for the largest adequately fitting arclets, but it will yield the same two values for extent (i.e., the number of arclets) and turning angle, which together specify the same scale-invariant contour shape.

Figure 3 illustrates this property. At bottom are two identical ellipses of differing sizes. For some portion of the boundary of the figure in (B), an approximately constant curvature segment is detected by arclets of a certain scale (k) having a certain turning angle (Θ). In the example, in both cases the constant curvature segments representing part of the two ellipses are formed from two cooperating arclets, encompassing three underlying oriented units, giving the value of this extent parameter n as 3 for this constant curvature segment. All three of these parameters are shown as the *scale specific code* in the diagram, which, for the smaller ellipse, is shown as $\{\Theta, k, n\}$. Turning to the larger ellipse in (A), we find that the largest, best-fitting arclet is at a different scale, 3k. At that scale, the constant curvature segment is fit by arclets that also have turning angle Θ and an extent of n units, shown as $\{\Theta, 3k, n\}$. Note that the scale-invariant code, which is identical for these two identical shapes, emerges simply from the turning angle parameter Θ and the number of units n needed to encode the constant curvature segment. This model is consistent with our effortless perception of identical shape for objects of different sizes (scale invariant code) and also our ability to see that these objects have differing sizes (scale-specific code).

This property of obtaining scale invariance for free comes from the use of oriented segments of finite lengths to encode curvature, along with encoding at multiple spatial scales. Consistent with what is known about the kinds of information encoded in early cortical processing, our ideas are built upon the assumption that extraction of curvature in visual processing comes from positions and relations of oriented units. Whereas the fact that the visual system uses oriented units of finite lengths to encode curves might seem to be a compromise or limitation in encoding what are actually continuous curves in the input, it is this characteristic, along with obtaining best fits in multiscale curve detection, that allows a scale-invariant curvature property to emerge automatically.

We have introduced this theoretical background at this point to illustrate important aspects of a constant curvature coding scheme in a multiscale framework. This framework provides a backdrop for all of the experiments reported here, and some of its more detailed consequences are further developed and tested in Experiment 3. The plan of the present work is as follows. We discuss briefly other research that bears on the possibility of constant curvature encoding of contour shape in human and primate vision. Then, we present the results of two experiments that furnish evidence from differing paradigms suggesting constant curvature representations in human vision. We then return to some detailed consequences of the arclets model, which motivate specific predictions that are then tested in Experiment 3.

Evidence Relating to Constant Curvature Encoding

At a general level, the arclet theory of contour shape representation is motivated by the idea that the informational cost of encoding a contour shape can be reduced, while adequately representing the stimulus, by encoding it in terms of constant curvature segments. These constant curvature segments are the "parts" of a structural description. The constant curvature model of shape representation is well suited to support invariance to planar transformations and is object-centric, describing a shape in terms of the relations among individual constant curvature segments. The resulting representations are translation, rotation, and scale invariant.

Ecological support for the idea that 2D projections of objects can be approximated by partial circles has been found in natural scene statistics, either because there are a large number of circular contours in natural scenes (Sigman et al., 2001), or because there are an abundance of smooth, closed contours in natural scenes that are well approximated by circles (Chow et al., 2002). There is also neurophysiological evidence suggesting that curvature at particular object-centered angular locations is coded in object related areas of primate visual cortex (Pasupathy & Connor, 2001) and that shapes may be represented as a collection of boundary fragments (Pasupathy & Connor, 2002).

Despite the theoretical support and neurophysiological evidence, there has been relatively little direct, behavioral evidence in support of the arclet theory up to this point. In one study, Garrigan and Kellman (2011) examined the visual system's accuracy in a contour matching task. Subjects were asked to decide if sequentially displayed open constant curvature and nonconstant curvature contours were the same or different shape, apart from a scaling transformation. They predicted that, if contour shape is encoded via constant curvature segments, performance for constant curvature stimuli should be better than for nonconstant curvature stimuli. The latter might require contour shape to be represented with lower fidelity or require more parts to be recoded in terms of constant curvature segments. The results showed that accuracy was

reliably higher when the open contour was made of segments of constant curvature than when it was made of segments of continuously changing curvature.

In the present work, we aim to further assess the empirical validity of constant curvature segments as primitives for shape representation in three new paradigms. In Experiment 1, we compared path detection (Field et al., 1993) for paths with constant turning angles and for paths with nonconstant turning angles. We predicted that if constant curvature has special perceptual significance, the paths with constant turning angle should be easier to pick out among randomly oriented distractors. In Experiment 2, we measured subjects' ability to segment contour fragments made of two segments of constant curvature. Segmentation performance in the constant curvature task was compared with performance on a control task in which the contour fragment was made up of two segments of constantly accelerating curvatures (Euler spirals). We hypothesized that if contour representations rely on constant curvature components but not changing curvature components, subjects should be more sensitive to changes in the curvature of a contour fragment than to changes in the rate of curvature. Finally, we tested a prediction of the constant curvature model: Changes to a contour's angular extent should be more detectable than changes in its absolute curvature. This is because while changes to absolute curvature can be explained by a difference in scale, changes to the angular extent of an arclet always indicate a different shape. Subjects performed a same/different task in which either absolute curvature or angular extent was changed and we measured subjects' ability to detect both kinds of changes to an open contour fragment.

Experiment 1

The arclet theory of contour shape representation has a number of consequences that may be observable in carefully designed psychophysical tasks. In Experiment 1, we tested the role of constant curvature in detection performance. We used a modified version of the path detection paradigm and predicted that constant curvature paths, presented among randomly orientated distractors, might be easier to detect than non-constant-curvature paths, because constant curvature paths have a simpler representation and/or greater perceptual unity than nonconstant curvature paths. In Experiment 1a, we compared search for constant curvature targets with search for targets with equal total curvature (i.e., the sum of unsigned curvatures between adjacent segments) but nonconstant curvature polarity. In Experiment 1b, we restricted both constant and nonconstant curvature paths to have a single curvature polarity while still equating total curvature.

In the path detection paradigm, an observer is typically asked to identify which of two sequentially presented arrays of variously oriented Gabor elements contains an embedded "path" (see Figure 4). A path can be defined with constraints on the relations of adjacent element pairs in many ways, but, strikingly, paths are much more easily detected when adjacent elements making up the path satisfy certain geometric relations (Field et al., 1993). These relations, described as an "association field" of linkages between oriented units (Field et al., 1993), appear to be identical to the geometry of contour relatability, which describes the spatial relations of edge fragments that produce contour interpolation in modal and amodal completion (Field et al., 1993; Hess & Field, 1999; Hess et al., 2000; Kellman et al., 2005; Kellman et al., 2001). A key difference is that paths do not typically give rise to perceived continuous edges between the inducing elements. Recent work suggests that the salience of paths reflects the operation of an intermediate contour-linking representation, where additional constraints must be fulfilled to produce perceived edges connecting path elements (Carrigan, 2020; Carrigan & Kellman, 2020; Kellman et al., 2016).

Figure 4
Illustration of the Path Detection Paradigm

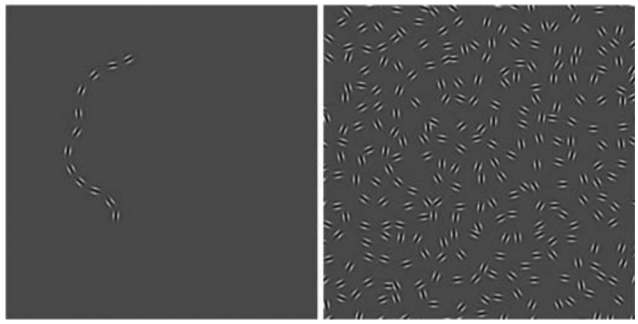

Note. In a typical experiment, two images containing randomly arranged Gabor elements are shown during each of two intervals. One of the images also contains a path. The participant's task is to report which of the two intervals contained the path. Left: Example of a path. Right: An array containing the path on the left. Certain relations of elements, as in the path on the left, facilitate detection of the path (From "Contour integration by the human visual system: Evidence for a local 'association field' " by D. J. Field, A. Hayes, and R. F. Hess, 1993, *Vision Research, 33*, pp. 173-193, Copyright Elsevier (1993)).

How does the present hypothesis of constant curvature encoding in contour shape representation relate to path detection and what is known about the conditions under which it occurs? Although we have described constant curvature encoding with respect to continuous contours given in the stimulus, the operation of arclets in that process might be involved in contour linking across gaps as well. It is interesting to note that some work in contour interpolation suggests such a relationship. Specifically, the geometric relationships that define the concept of relatability have a natural relationship to a particular form of interpolated edges: Every relatable edge can be described as a combination of a constant curvature segment and a zero curvature segment (see J. Skeath, Appendix B, in Kellman & Shipley, 1991). Earlier, Ullman (1976) suggested that every illusory contour is the combination of two constant curvature segments. If, as we suggest in this work, contour shape is encoded as constant curvature segments, the same machinery that encodes contour shape may be involved in contour interpolation across gaps.

Descriptions of the geometric relations involved in the association field or contour relatability have not typically suggested a role for constant curvature (but see Pettet, 1999). Except for the property of closure, geometric properties relevant to path detection have been definable between members of each pair of elements. For example, angular deviation from collinearity has been found to decrease the strength of path detection (and contour interpolation). Constant curvature encoding suggests that some relations that encompass more than one element pair may be important. In

Experiments 1a and 1b, we tested whether constant curvature relations matter in path detection, where other known influences, such as the average angular relations between elements, were controlled.

Experiment 1a

Experiment 1a adapted the path detection paradigm developed by Field et al. (1993), in which paths were defined by the pairwise relations of orientation and positions of a set of Gabor elements, embedded in a field of randomly oriented Gabors. On every trial, observers viewed two screens sequentially in a two-interval forced choice (2IFC) procedure. One of the displays contained distractor elements and the target path; the other contained only distractor elements. Field et al. (1993) found that participants could more easily detect the path if the Gabor elements satisfied certain positional relations, and performance also decreased as the orientation difference between adjacent elements along the path increased.

Experiment 1a differed from the design of Field et al. (1993) in a number of important ways. First, instead of Gabor elements, the targets and distractors were composed of line segments. Line segments were used instead of Gabor elements because we are specifically interested in relations of contour elements. Second, contour elements in the display were allowed to overlap, while the Gabor elements in the original paradigm were constrained to grid locations such that they could never overlap. Distractor segments were positioned randomly (and could therefore occasionally overlap) because this experiment required more flexibility in the path shapes than in the studies of Field et al. (1993). Constraining the target elements to lie on a grid also constrains the types of shapes that can be created from the elements. Last, we used a present/absent detection task instead of the two-interval forced choice paradigm. A present/absent paradigm was used instead of a 2IFC design to discourage strategies not related to detection of the shape that might have more to do with the statistical properties of the geometric configuration of the target and distractor elements as a whole.

Method

Participants

Initial data estimated an effect size $\eta_p^2 > 0.5$ for constant versus nonconstant turning angle, so we estimated that we should have at least 16 participants for a power of .8. Participants were 17 undergraduates from the University of California, Los Angeles who received course credit for participation. All participants reported normal or corrected-to-normal vision. No participants were excluded. All procedures completed by participants in this study were approved by the University of California, Los Angeles (UCLA) Institutional Review Board.

Stimuli

The stimuli were of two types: a target hidden in noise and noise alone. The target consisted of six white line segments 17.6 arcmin long and 4.8 arcmin wide, each oriented at angle θ relative to the last segment and separated by 13.2 arcmin along the path. θ had values of $\pi/9 + 2^{*}\pi^{*}n/45$ with n taking integer values between 1 and 5 (corresponding to five angular relations: 28, 36, 44, 52, and 60 degrees). These curvature values appeared to be easily discriminable and created a set of targets spanning from low (but nonzero) curvature to a fully closed curve. The noise consisted of white line segments randomly oriented and positioned within a 13° × 13° square with uniform density of approximately 1.5 line segments per deg^2. Consequently, there were, on average, 253.5 distractor line segments distributed uniformly across the square region upon which the target was superimposed. All line segments were presented on a uniform gray background.

Targets were constrained to lie within a circular region 5.6° in circumference, positioned at the center of the independently generated field of distractor segments. On "target present" trials, the target was superimposed on the distractor line segments, within this region. Since distractor segments could overlap other distractor segments, the potential overlap of target segments and distractor segments did not reliably signal that the target was present. On "target absent" trials, six additional distractor elements were included, randomly oriented and positioned within the same circular region that constrained the targets. This ensured that the average number of segments (target + distractors) was equated between target present and target absent conditions. Similarly, the density of segments, and the small radial dependency of segment density, was closely matched across conditions.

Targets had curvature of two types, same polarity (SP) and alternating polarity (AP). Locally (between two adjacent segments), the SP and AP targets were indistinguishable. In the case of the SP targets, the angular change between any two adjacent segments was always in the same direction (always turning to the right or always turning to the left). In the case of the AP targets, the direction of the angular turn between any two segments was always opposite to the preceding pair of segments. For any pair of elements, then, the difference in orientation on an SP trial would be the same as the difference in orientation on an AP trial with matched turning angle. All targets are shown in Figure 5.

This experiment is similar to earlier research on long-range interactions in contour detection (Pettet, 1999). The stimuli used in the earlier work were more similar to those used by Field et al. (1993) in that the targets and noise were composed of oriented Gabor elements that were constrained to not overlap spatially. Pettet (1999) maintained a constant turn-angle between local elements and varied how frequently the path changed direction. There was only one condition in which the path changed direction at the highest frequency (every third element, immediately after the path direction is established). As Pettet's (1999) experiment relates to the current research, this was the most important condition because any lower frequency of path-direction change allows for detection of the path by detecting an SP subpath within the stimulus. For this reason we prefer our design, which is similar to the highest-frequency path-direction change condition in Pettet (1999).

Design and Procedure

We used a within-subjects design, testing participants on trials with both AP and SP targets. Participants were also exposed to both SP and AP targets during training. The experimental session consisted of 260 trials.

Figure 5
Targets for Experiment 1a

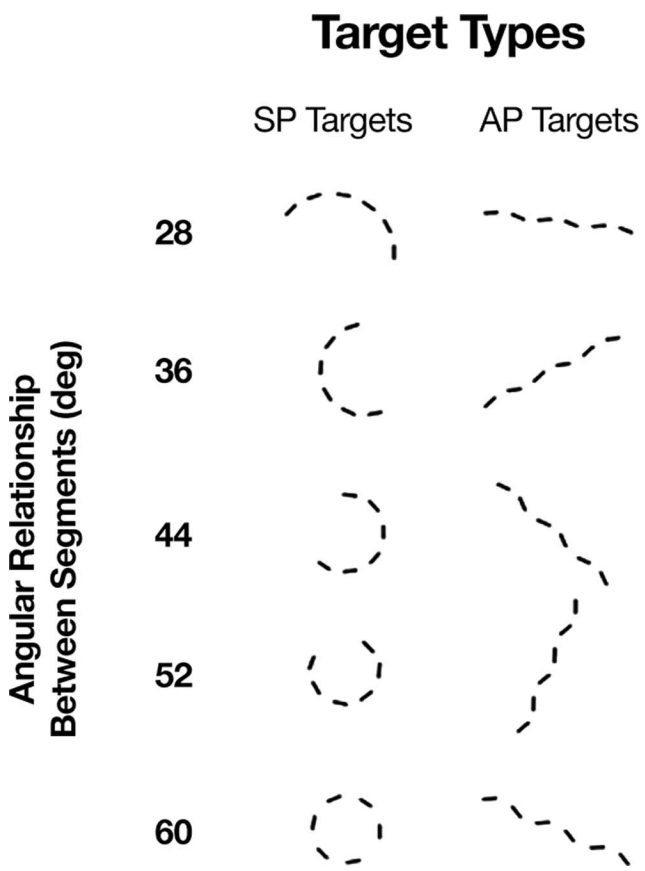

Note. Same polarity (SP) targets (left) and alternating polarity (AP) targets (right) are shown for all angular relation values.

In this experiment, and all experiments presented in this article, participants were seated in a chair with their head stabilized in a chin rest 78 cm in front of a 20-in. ViewSonicTM P225f Monitor with 1,152 × 870 resolution (approximately one pixel per 0.35 mm^2), operating at 75 Hz refresh rate. On each trial, a fixation point (a plus sign) was presented at the center of the display. After 750 ms, the fixation point was replaced with either a target-present or a target-absent stimulus. On the target-present half of trials, half had an SP target and half had an AP target. The target types were randomly interleaved throughout the experiment. After a 300-ms presentation time, the screen was replaced with a mask (see Figure 6). Participants were required to make a keyboard response indicating whether a target was present or not. After a keypress was made, feedback was given using two distinct tones for correct and incorrect responses and the next trial began. Prior to training, participants were familiarized with both the experimental task and the target stimuli through 40 practice trials. Practice trials were identical to standard test trials, except that the noise density gradually increased throughout the practice session from .49 segments per deg^2 to 1.5 segments per deg^2.

Results

Results from Experiment 1a are plotted in Figure 7A, which shows sensitivity (d') as a function of different angular relations between adjacent line segments. Hits were scored as correct detections of a path, and false alarms were scored as incorrect reports of a path in trials on which no path was presented. The false alarm rates were therefore the same in both the SP and AP condition, and differences in sensitivity were solely driven by hit rates. The average false alarm rate across subjects was 13%. Effect sizes for the ANOVA were calculated using an adjusted measure of partial eta squared to reduce positive bias (Mordkoff, 2019). Performance for SP paths was consistently better than for AP paths. This effect was confirmed by statistical analyses. A 2 (condition) × 5 (angular relation) repeated measures analysis of variance, showed a reliable main effect of condition, $F(1, 16) = 88.64$, $p < .01$, adjusted $\eta_p^2 = .84$ and angular relation, $F(4, 64) = 2.93$, $p = .028$, adjusted $\eta_p^2 = .11$. There was no reliable interaction of stimulus class and angular relation, $F(4, 64) = 0.73$, $p = .55$, adjusted $\eta_p^2 < .001$. In Figure 7B, the data are replotted after removal of seven participants who performed near ceiling in the SP condition. Some participants performed near ceiling (d-prime >3.0) in the SP condition, but no participants performed near ceiling in the AP condition. Some decrease in sensitivity with increasing angular relation between adjacent elements is evident, as well as increased sensitivity at the highest angular change as the contour forms a closed loop. The decrease in sensitivity with increasing angular relation is consis-

Figure 6
Schematic of a Trial in Experiment 1a

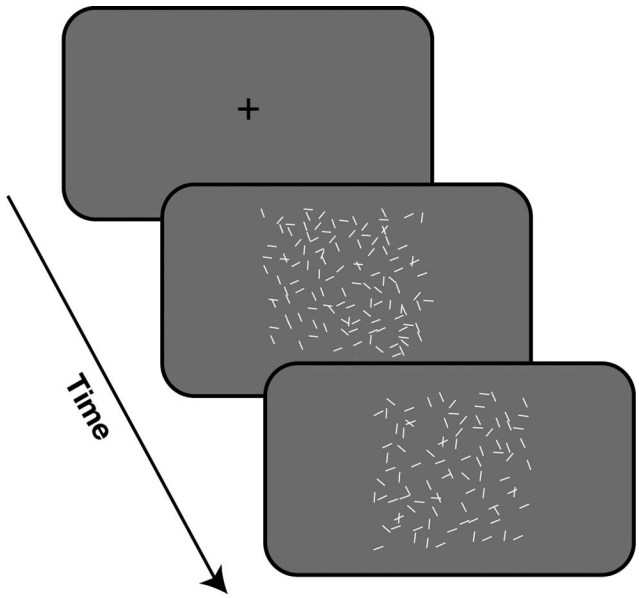

Note. Participants first viewed a fixation point, which was then replaced by a target hidden in noise or noise alone. Targets were composed of six white contour elements arranged in one of the target configurations. The noise was a field of randomly oriented line segments like those that formed the target. After a short period of time, a mask, statistically identical to the noise, was presented and remained until participants made a keyboard response indicating the presence or absence of a target on the preceding screen. In the trial shown here, an SP target is present.

Figure 7
Results From Experiment 1a

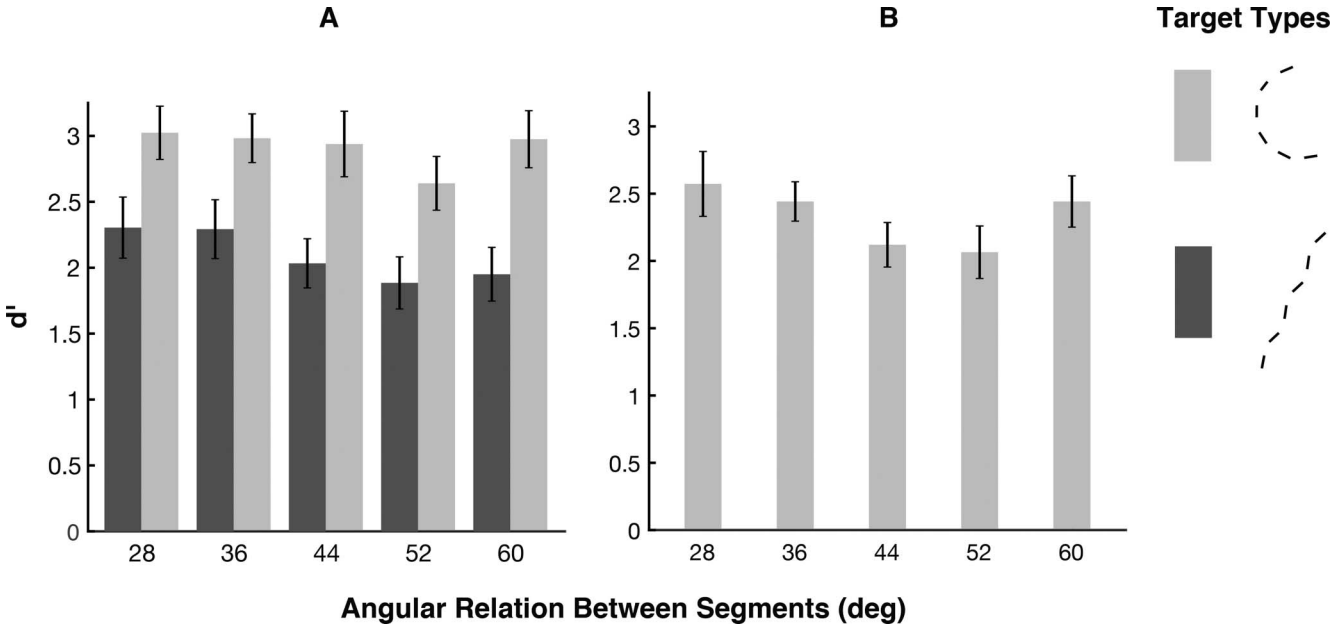

Note. Sensitivity is plotted for detection of targets with different angular relations (deg) between path segments. (A) Participants showed greater sensitivity for detection of SP targets relative to AP targets. (B) SP target detection sensitivity is plotted with participants with sensitivity near ceiling ($d' > 3.0$) excluded. This criterion was never reached in the AP condition. A decline in sensitivity with increasing angular relation between target segments is observed, as well as an increase in sensitivity due to contour closure (for targets with an angular relation of 60° between adjacent elements).

tent with the results of Field et al. (1993). The closure effect is consistent with Kovacs and Julesz (1993), who found that the maximum interelement spacing for a detectable contour was significantly higher for closed contours.

Discussion

In this experiment, participants were presented with paths of two types, one with constant curvature polarity, (SP), such that all orientation changes were in the same direction, and one with alternating curvature polarity, (AP), such that the direction of orientation changes between adjacent segments alternated from clockwise to counterclockwise. Participants were reliably better at detecting the contours in displays with an SP target than in displays with an AP target. Subjects' better performance on trials with an SP target cannot be explained by higher overall curvature in the SP condition, as total curvature was equated between the AP and SP conditions. The only difference between SP contours and AP contours lay in the variation in turn direction between adjacent element pairs. While the SP target had consistent turn direction, the direction in the AP target reversed between adjacent pairs.

The results of Experiment 1A are consistent with the hypothesis that contour segments are encoded into components of constant curvature. This result could arise because oriented units are linked into higher-order curvature detectors (arclets) that are activated more strongly by more segments having the same curvature than occurs with paths made of nonconstant curvature elements. A variant of this hypothesis is that paths of constant curvature segments produce a simpler representation (a constant curvature segment) than paths of nonconstant curvature.

The geometric properties that determine the strength of path detection (Field et al., 1993) are consistent with the geometry of *contour relatability* that determines which edges give rise to contour interpolation (Hess & Field, 1999; Kellman & Shipley, 1991; cf., Parent & Zucker, 1989). The differences between SP and AP in Experiment 1, however, cannot be attributed to the geometric relations between pairs of elements. Rather, they suggest an additional factor: Constant curvature, established over at least three spatially separate units, affects detection sensitivity. These results provide evidence that constant curvature detectors and/or the use of constant curvature in contour representations enhances detectability independent of the pairwise relations known to be important in path detection and contour interpolation.

There are two limitations of Experiment 1a. First, the polarity of curvature is confounded with constant and nonconstant curvature, and so we cannot be sure that the advantage for the constant curvature targets in Experiment 1a was not due to a consistent turn direction between adjacent elements, rather than a consistent turning angle between adjacent elements. Second, a detection task, like the one used, does not guarantee, or even necessarily encourage, participants to encode all parts of the target. Shape is a relational property, defined, in our case, over entire contours. If part of a shape is substantially easier to detect than another part of a shape, then an obvious strategy in a detection task is to search not for the entire shape, but just for the part that is easy to detect. To study performance related to shape encoding, we needed a task that

required encoding the full target in order to attain high levels of performance.

It is for these reasons that two design changes were implemented in Experiment 1b. First, all stimuli were composed from elements with the same curvature polarity, but with either high, low, or mixed (part high, part low) curvature. We expected mixed (i.e., nonconstant) curvature, stimuli to be harder to detect than both high and low (constant) curvature stimuli even though the local pairwise geometry between neighboring elements would suggest these stimuli should be easier to detect than stimuli composed from high curvature elements, and harder to detect than stimuli composed from low curvature elements. Second, to measure sensitivity to contour shape, it is important to be sure task performance depends on detecting all parts of the stimulus. A detection task, like the one used in Experiment 1a can be performed through partial detection, and so performance differences may be driven by the local geometry of parts of the stimulus. Similarly, partial detection could account for another result in Pettet (1999), in which detection of a nonconstant curvature stimulus composed of pairs of elements, each taken from a constant-curvature stimulus, was similar to the average of all the constant curvature stimuli from which its neighboring element pairs were taken. In our Experiment 1b, we employed a contour shape-matching task, using four contour shapes designed so that the angular relation between any two neighboring elements from any shape could come from three of the four candidate shapes. Consequently, detecting even three of four neighboring elements was insufficient for identifying a shape, and so shape matching required detection of all parts of both shapes.

Experiment 1b

The results of Experiment 1a are consistent with the hypothesis that the target contours are represented in terms of constant curvature units and that search is therefore easier when targets can be represented with a smaller number of longer constant curvature segments. However, as Pettet (1999) pointed out, such results are also consistent with a preference for grouping paths with constant curvature polarity but not necessarily the same curvature. Pettet tested this hypothesis by comparing detection performance for six constant curvature paths to detection performance for a constant curvature polarity path composed of elements related by one turn-angle from each of the six constant curvature paths, configured in order of curvature magnitude. A similar target, using four line segments instead of seven Gabor elements is shown in Figure 8. Pettet's (1999) results showed no significant difference between the average of the six paths and the hybrid nonconstant curvature path, suggesting that there was no preference for cocircular grouping in paths with constant curvature polarity.

These results, however, do not control for the local angular relations among pairs of elements. Specifically, the constant curvature condition averages across low-curvature, easily integrated paths and high-curvature, less easily integrated paths. In Experiment 1b, we aimed to distinguish whether sensitivity to integration or constant turning angle played a larger role in contour detection. If integration sensitivity is the more important factor, we expected constant curvature paths with high overall curvature to be more difficult to detect than nonconstant curvature paths with lower overall curvature. On the other hand, if constant curvature was the more important factor, we expected paths with constant turning angle to be easier to detect than nonconstant turning angle paths, even if their overall curvature was higher.

Figure 8
Nonconstant Curvature Targets, as in Pettet (1999)

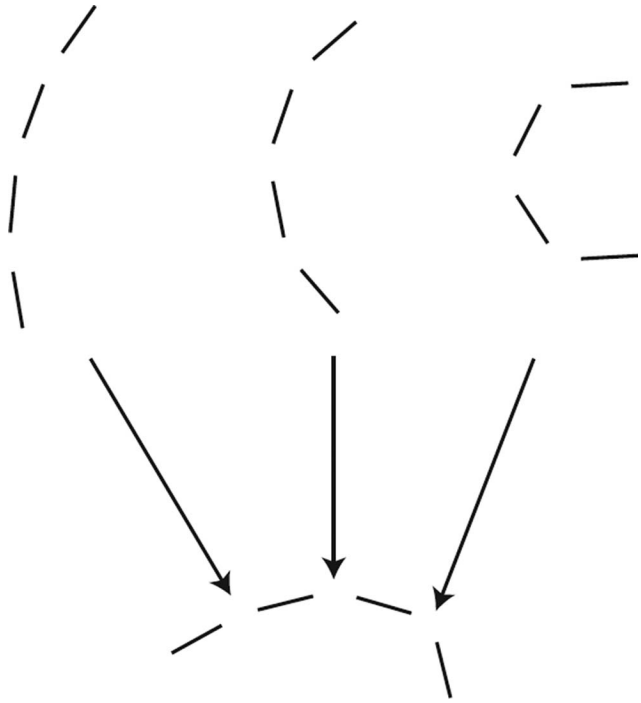

Note. Three constant curvature targets are shown (top). The bottom figure is composed from element pairs related by one of the turning angles from each of the three figures above it. These targets are similar to those used by Pettet (1999), except that they are composed of four line segments instead of seven Gabor elements.

Method

Participants

Initial data predicted that we would have an effect size of Cohen's $d > 1$ for both the low and high curvature conditions versus the mixed condition, so we would need at least 11 participants. Sixteen additional undergraduates from UCLA, who were not participants in Experiment 1A, were recruited and received course credit for participation. All participants reported normal or corrected-to-normal vision. No participants were excluded. All procedures completed by participants in this study were approved by the UCLA Institutional Review Board.

Stimuli

The stimuli in Experiment 1b were similar to the stimuli from Experiment 1a. Each consisted of a target presented near the center of a field of randomly oriented and positioned noise elements with noise density of 0.48 segments per deg^2. The targets, however, were of four types. Two constant curvature targets were defined by four line segments, each 17.6 arcmin long and evenly spaced with 13.2 arcmin gaps. In one target, the turning angle between seg-

ments was 12°, while in the other the elements related to each other by a turning angle of 48°. There were also two nonconstant curvature targets, which we refer to as "mixed" curvature. They were made up of four line segments with the same length and spacing as in the constant curvature conditions. The mixed targets shared three angular relations. In one target, the first two elements related by a turning angle of 12° and the last related by a turning angle of 48°. The other had the opposite configuration of element relations. All four configurations are shown in Figure 9.

During the experimental session, targets were always presented in a field of randomly positioned and oriented line segments 17.6 arcmin long. These line segments, as well as the targets were white, presented on a uniform gray background. An additional noise mask of randomly oriented and positioned black line segments of the same dimensions was also generated on each trial. This noise mask had no target concealed in it. The noise mask elements were black to distinguish the masks from the stimuli.

Design and Procedure

Participants were required to view two consecutively presented stimuli and determine if the shapes concealed in the noise of both stimuli were the same or different. The aim of this paradigm was to assess differences in path detectability while also ensuring that the full path had to be detected for responding to be reliably accurate. Participants would need to detect the whole path, not just a few elements in a path, in order to accurately perform the same/different task. On each trial, a fixation point (a plus sign) was presented at the center of the screen for 200 ms. After 200 ms, the fixation point was replaced by the first stimulus which remained visible for 200 ms. This stimulus was then replaced with the same fixation point which again remained visible for 200 ms. The second stimulus then replaced the fixation point and remained visible for 200 ms, followed by the noise mask which remained visible until a keypress was made (see Figure 10). Participants were instructed to make one keypress if they thought the two shapes in each of the stimuli were the same, and a different keypress if they thought the two shapes were different. After a keypress was made, feedback was given using two distinct tones for correct and incorrect responses, and the next trial began.

There were 80 practice trials and 240 experimental trials. The practice trials were the same as the experimental trials except that the noise density increased from zero (no noise elements) to 0.48 segments per deg^2 during the course of the practice session. This allowed participants to learn to perform the task and familiarize themselves with the potential targets. During the experimental session, participants were given a 60 second break every 60 trials.

We once again used a within-subjects design for Experiment 1b. On half the trials, same shapes were presented, on the other half, different shapes were presented. High, low, and mixed curvature trials appeared equally often on same and different-shape trials, ensuring that detecting only one of the stimuli did not benefit performance. All of the different trials included a presentation of one of the mixed stimuli and one of the constant curvature stimuli, with each type of stimulus presented equally often. Because the mixed stimuli consisted, on average, of equal parts high and low curvature and mixed stimuli were presented as often as high or low curvature stimuli, detection of part of a stimulus could not be used to reliably identify the entire stimulus. Target subparts, which were either high or low curvature in all conditions, were equally likely to appear in the context of a same or different trial.

Figure 9
Targets for Experiment 1b

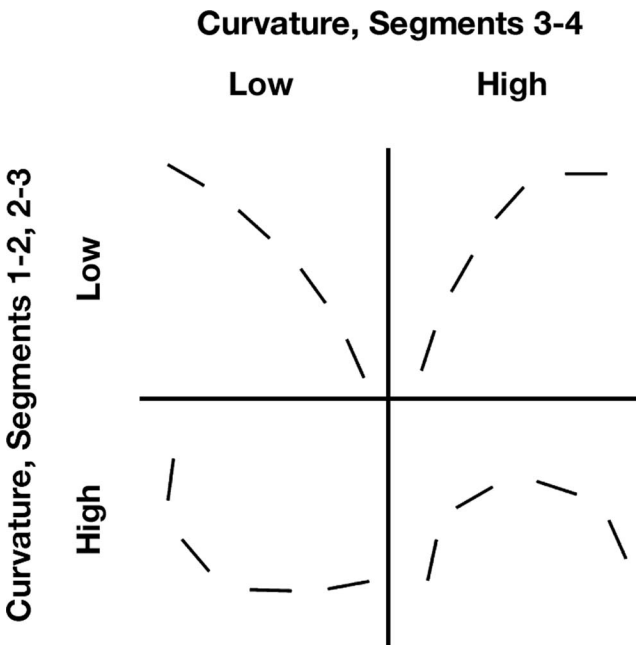

Note. Targets were composed of four line segments, which could be arranged in any of the configurations shown. Configurations had either three low turn-angle relations, three high turn-angle relations, or a mix of low and high turn-angle relations between target elements.

Results

Results from Experiment 1b are plotted in Figure 11. For each stimulus type, sensitivity was calculated by comparing performance on trials on which the same shape was presented twice to all trials on which different shapes were presented. Differences in sensitivity across target types were therefore entirely determined by trials on which only one shape type was presented. The false alarm rate across subjects was 35%. Effect size was calculated as Cohen's d_{av}, a measure of effect size with an adjustment for within-subjects comparisons that averages the standard deviation from both conditions (Cumming, 2012; Lakens, 2013). There was a significant effect of target type, $F(2, 30) = 22.244$, $p < .01$, adjusted $\eta_p^2 = .570$. Paired t tests showed no reliable difference in sensitivity between the two constant curvature conditions, $t(15) = 1.033$, $p = .318$, Cohen's $d_{av} = 0.30$. There was a reliable difference between the low curvature CC condition and the mixed curvature condition, $t(15) = 5.401$, $p < .01$, Cohen's $d_{av} = 2.02$ and between the high curvature CC condition and the mixed curvature condition, $t(15) = 6.409$, $p < .01$, Cohen's $d_{av} = 1.44$.

We also examined sensitivity by comparing performance on "same" trials for each stimulus type to only those different trials on which the same type of shape was shown in one of the displays. For example, sensitivity for high curvature contours can be calculated as subjects' hit rate when shown two high curvature contours versus their false alarm rate when shown a high curvature contour

Figure 10
Paradigm for Experiment 1b

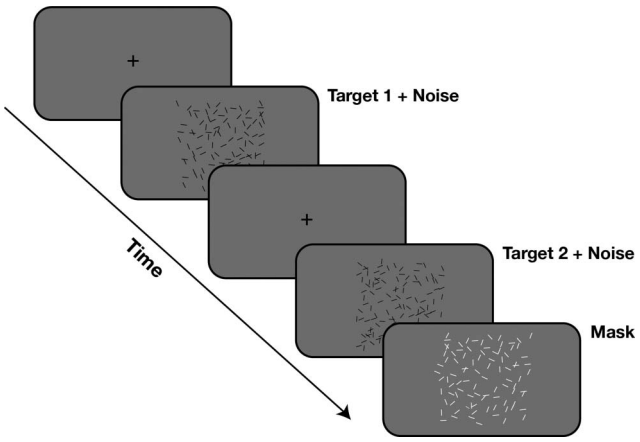

Note. In this experiment, participants were required to find two targets concealed in noise and compare their shapes. On each trial a fixation point would first be presented, followed by a target concealed in noise, a second fixation point (identical to the first), a second target concealed in noise, and finally a mask of noise alone. Participants would make a keyboard response indicating if they thought the targets had the same or different shapes.

and a mixed contour. This more detailed analysis revealed similar effects to those reported above. Using this calculation, false alarm rates were 30% for high curvature contours, 39% for low curvature contours, and 37% for mixed curvature contours. In the analysis of sensitivity using these false alarm rates, there was a significant effect of target type, $F(2, 30) = 20.80$, $p < .01$, adjusted $\eta_p^2 = .55$. Paired t tests showed no reliable difference between the two constant curvature conditions, $t(15) = 1.31$, $p = .21$, Cohen's $d_{av} = 0.43$, but did show reliable differences between the low curvature CC condition and the mixed curvature condition, $t(15) = 4.91$, $p < .01$, Cohen's $d_{av} = 1.66$ and between the high curvature CC condition and the mixed curvature condition, $t(15) = 6.60$, $p < .01$, Cohen's $d_{av} = 1.84$.

Discussion

Paths with constant turning angle were easier to detect among distractors than paths with varied turning angle, regardless of whether the mean turning angle of the constant curvature stimuli was higher or lower than in the mixed condition. These results indicate that the effect observed in Experiment 1a was at least partly due to the constant curvature of some of the stimuli, and not merely their consistent curvature polarity. As in Experiment 1a, the results of Experiment 1b can be explained by the visual system encoding integrated contours as a set of constant curvature arclets. The mixed curvature targets cannot be accurately represented by a single segment of constant curvature, while the paths with constant turning angle can, and are therefore representationally simpler.

The results of Experiment 1b would not be predicted by other approaches to contour shape representation. For example, Attneave (1954) theorized that, from the perspective of information theory, the least surprising continuation of a contour was a straight line sharing the tangent direction of the terminating point. Feldman and Singh (2005) extended this idea to develop a metric for contour complexity in which contours with higher curvature had more bits of information than contours of low curvature. Other theories from computer vision make similar predictions. For example, Davis (1977) developed an algorithm in which shape is represented by a set of straight lines, with breaks recursively added to the description until all straight-line approximations are below a certain threshold. Contrary to predictions of these theories, constant curvature paths were simpler than mixed paths regardless of whether their average curvature was higher or lower than the average curvature of the mixed paths, suggesting that curvature regularity, not curvature itself, is a better predictor of the representational complexity of a contour.

A different question is how the current results relate to theories that aim to describe the natural parts of objects. Hoffman and Richards (1983) theorized that parts are perceived between local minima. If detectability depends on the number of parts, one might apply this idea to Experiment 1a, where the displays with alternating polarity stimuli have many local minima, which might explain differences in performance between them and constant curvature paths. However, all stimuli in Experiment 1b are monotonic, and would be coded as a single part in Hoffman and Richards's (1983) theory. More broadly, the advantages in detection we found in Experiments 1a and 1b may involve an earlier perceptual process of representing contour tokens and their shapes. Salient perceptual parts may involve relevant ecological and functional issues in object perception that arise subsequent to a more basic level of description in visual perception that involves initial representation of contours. The visual system appears to encode monotonic contour segments with several primitive elements based on variation in curvature.

Figure 11
Results From Experiment 1b

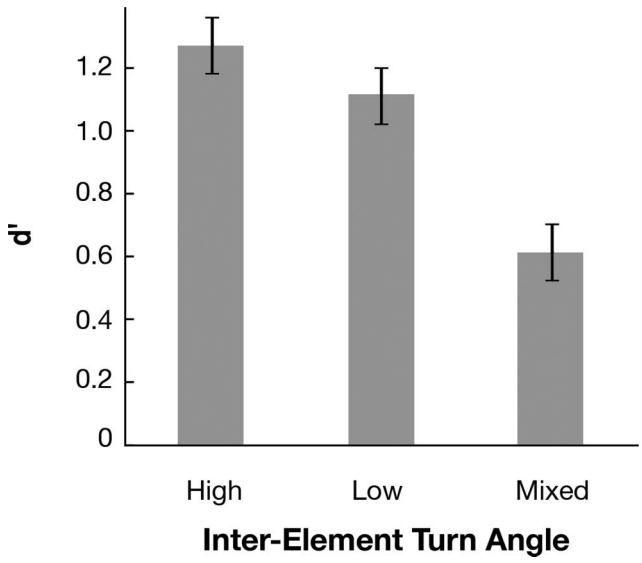

Note. Participants had higher sensitivity for comparing high and low curvature targets composed of constant curvature segments, compared with the mixed targets condition. Error bars are ±1 standard error of the mean.

Experiment 2

A representation formed from constant curvature parts has favorable attributes (e.g., translational, rotational, and scale invariance) and neurophysiological plausibility. Experiment 1 provides evidence that contour shapes composed of regions of constant curvature are easier to detect and/or represent. These results are consistent with a representation of contour shape composed of joined constant curvature segments. An alternative view, however, is that the visual system may exploit any discoverable regularity. In Experiment 2, we test this theory using the next higher-order primitive beyond curvature, defined as regions of constant change in curvature per unit of arclength. A contour with a constant first derivative of curvature per unit arclength, known as the Euler spiral, is expressed by the parametric equations:

$$x = \alpha \int_0^t \cos u^2 du = \alpha C(t)$$

$$y = \alpha \int_0^t \sin u^2 du = \alpha S(t)$$

An example is shown in Figure 12. In the arclet theory, the curvature (change in orientation per unit arclength) of the contour is considered, and primitives are formed by locating regions of constant curvature or approximately constant curvature. The change in *curvature* per unit arclength (referred to here as the "acceleration" of the curve) can be used in a similar way, forming primitives defined as regions of constant acceleration (i.e., Euler spirals). Research on the shape of subjective contours has previously proposed that the interpolating edge between two tangent discontinuities has the form of an Euler spiral (Kimia et al., 2003).

Does the visual system code contour shape with higher order curvature relations, or is curvature the basic unit of encoding? We test this in Experiment 2 using a new paradigm to compare the visual system's sensitivity to changes in acceleration with its sensitivity to changes in curvature. This experiment is based on the assumption that if the visual system is sensitive to some quantity on a curve, observers should be able to locate positions at which that quantity changes, provided the magnitude of the change is large enough. We hypothesized that if constant curvature is the basic unit of encoding, participants should be able to learn to segment a smooth contour into two constant curvature segments but not into two segments with different constant curvature accelerations.

There are two reasons for using a different paradigm in Experiment 2 from the ones used in Experiment 1. First, establishing a path with a higher-order relationship among elements requires more segments, overly constraining types of paths that can be constructed. The second reason is more connected to the overall purpose of the present work: If constant curvature components are fundamental to contour detection and encoding, we should be able to find evidence of them in a variety of tasks involving the encoding, recall, or comparison of contour shapes.

Figure 12
Example of an Euler or Cornu Spiral

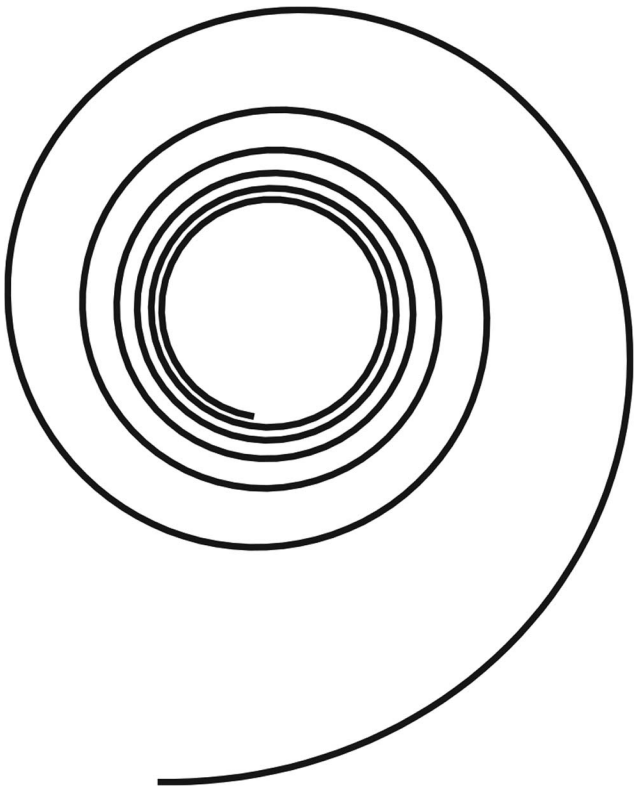

Note. In an Euler, also called Cornu, spiral, curvature changes linearly with distance along the contour.

Method

Participants

Pilot results suggested most (80%) participants could learn to divide the contour made from constant curvature segments, but none could learn to divide the contour made from Euler spirals. Because the pilot participants were experienced psychophysical observers, we used a more conservative 60% estimate. Power analyses therefore suggested we should have eight participants per condition for 80% power. Participants were 36 undergraduates from UCLA who received course credit for participation. All participants reported normal or corrected-to-normal vision. No participants were excluded. All procedures completed by participants in this study were approved by the UCLA Institutional Review Board.

Stimuli

Stimuli were of two types: piecewise constant curvature stimuli and piecewise Euler spirals (formed from two segments having different constant accelerations). Each contour was composed of two regions whose slopes matched at the join point. Each segment was constrained to pass through between 90° and 135°. For each stimulus type, the magnitudes of the relevant quantity (curvature or acceleration) for the individual regions along the contour were 2.5 times greater in one segment than in the other. All stimuli were uniformly scaled so that their longest extent along the horizontal or vertical axis did not exceed 7.6° of visual angle.

The average curvature difference between the segments of the constant curvature stimuli was equated to the average curvature difference between the segments that formed the constant acceleration stimuli. Constant acceleration stimuli were constructed so

that the smallest acceleration was subjectively detectable to the experimenter.

Design and Procedure

Each participant was tested in one of two conditions. Half the participants were presented with stimuli composed of two smoothly joined constant curvature segments. The other half were presented with stimuli composed of two smoothly joined constant acceleration segments. The participants' task was to segment each stimulus into two parts by indicating a point along the contour.

We did not assume that accurate segmentation would occur immediately; in fact, we expected that it would emerge across a number of learning trials. The reason is that a contour that is smooth (well-defined tangent or first-derivative) at all (nonterminal) points tends to be perceived as a physically connected entity (Kellman et al., 2003; Wertheimer, 1923). The constant curvature components that we theorize may underlie contour shape representation are intended to provide a shape description for continuous contours, not to indicate segmentation boundaries in the perception of separate objects in the world. Thus, the task assumes that with practice learners may come to access the underlying shape representation components, if they exist. Conversely, if such components do not exist, there will be little improvement with practice.

On each trial, a white and black contour stimulus appeared on a gray background, with a segmentation point randomly placed along its extent. The position of segmentation was designated by the point at which the contour changed color from white to black. Participants were instructed to adjust the position of segmentation, using a programmable joystick, to the position on the contour where it most naturally segmented into two pieces. When the participants were satisfied with their selected segmentation position, they were instructed to press a button on the programmable joystick. After pressing the button, the stimulus would be replaced by the same contour with two positions marked on it. One position corresponded to the segmentation proposed by the participant. The other position corresponded to the segmentation of regions by constant curvature or constant acceleration, depending on the condition of the experiment. Participants were never told any rule by which the "correct" target segmentation point was determined; it was their task to discover on their own how to properly segment the contour. A number corresponding to the distance along the curve between the two marks (in units of length equal to 5 arcmin, or one increment of the adjustable joystick), was also displayed. A diagram illustrating a typical trial for a piecewise constant curvature stimulus is shown in Figure 13, and three sample stimuli for the Euler spiral and constant curvature conditions are shown in Figure 14.

Participants were told the experiment would continue for 256 trials, or until they reached the performance criterion of 10 consecutive trials where the distance between the participants' segmentation and the segmentation into regions of constant curvature or constant acceleration was less than 10 units (50 arcmin).

Results

In the constant curvature condition, 12 out of 18 participants reached the learning criterion. On average, it took participants in the constant curvature condition 123 trials to learn to criterion. No participants reached the learning criterion in the constant acceleration condition. The average number of trials completed for each condition is shown in Figure 15. We tested for independence between reaching criterion and the segment condition (constant curvature vs. constant acceleration) with Yates's correction and found a significant association between segment condition and learning, $\chi^2(1) = 10.08$, $p < .01$.

Figure 16 shows, for each participant, mean error (in arcmin) across all responses, as well as the number of trials completed prior to reaching the learning criterion. Participants who did not reach the learning criterion (DNRC) completed 256 trials. Although the learning criterion was essentially arbitrary and very demanding (10 consecutive trials within 50 arcmin of the theoretical join point), most participants in the constant curvature condition did reach criterion, with higher mean error rates associated with more trials completed before reaching criterion. Even among participants who did not reach criterion, participants in the constant curvature condition had lower mean error, indicating these participants were also better able to learn to segment the contours, as compared with participants in the constant acceleration condition.

We also analyzed participants' pattern of errors to test whether they tended to place the transition point in the higher curvature or lower curvature contour region. For the constant curvature condition, participants on average placed the transition point in the low and high curvature region with about equal frequency (48% vs. 52%), $t(17) = .47$, $p = .64$. Participants' average signed distance from the transition point was 0.85 arcmin toward the high curvature region, which did not significantly differ from 0, $t(17) = -0.29$, $p = .78$, indicating little systematic bias. For the Euler spiral condition, participants were slightly biased to place the transition point in the high-acceleration region (58% of responses), $t(17) = -4.7$, $p < .01$. Their average signed distance from the transition point was also slightly biased toward the high-acceleration region (12.1 arcmin), $t(17) = -5.0$, $p < .01$. This

Figure 13
Paradigm for Experiment 2

Note. On each trial, participants were required to segment a contour formed from two constant curvature segments, or two constant acceleration segments. The task was to learn to segment the contour between the two constant curvature segments or between the two constant acceleration segments. After segmenting the figure, participants were shown their segmentation, the preferred segmentation, and a score indicating the distance between the two segmentations. Note that the examples of segmentation by color in the figure use black and red colors, whereas in the actual experiment, the contour segments were given in white and black. See the online article for the color version of this figure.

Figure 14
Sample Stimuli for Experiment 2

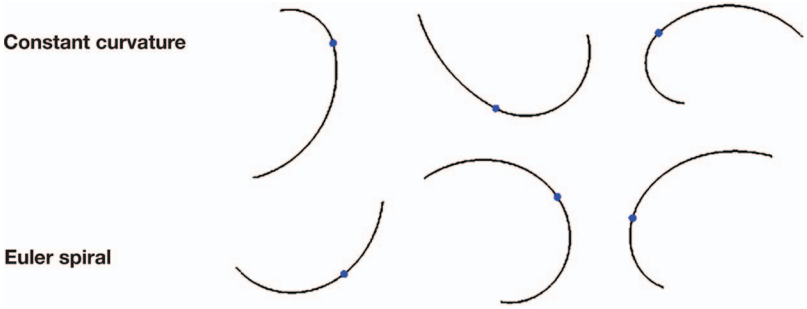

Note. Top: Contours made up of two constant curvature segments. Bottom: Contours made up of two Euler spirals. The blue dot on each contour indicates the correct transition point. See the online article for the color version of this figure.

distance is modest relative to participants' average error magnitudes (see Figure 16).

Discussion

The purpose of Experiment 2 was to find evidence of structural descriptions in contour shape representations by testing whether participants could learn to access hypothesized components in a contour segmentation task. We tested participants' ability to locate boundaries between joined constant curvature and joined constant acceleration segments. We found that participants could learn to segment contours into constant curvature parts with a small number of feedback trials, but could not learn to segment contours into constant acceleration parts. The learning in the constant curvature condition suggests that participants can access the underlying components of a constant curvature-based contour shape representation.

We believe different levels of performance at the conclusion of training does not simply reflect differing amounts of training required to segment the two different types of stimuli. In the constant curvature condition, learners showed increasingly accurate and precise responses from the beginning of training and typically reached a demanding performance criterion long before reaching the maximum number of trials. This kind of rapid acquisition and precision in performance seems more consistent with learning to access and use an existing representation for a novel purpose. The result stands in stark contrast to the Euler spiral segmentation condition, where participants floundered, showing

Figure 15
Comparison of the Number of Trials for the Constant Curvature and Constant Acceleration Conditions

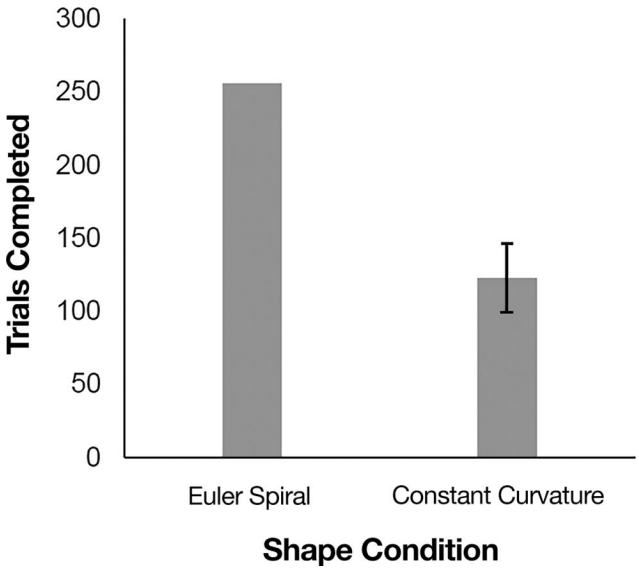

Note. For each participant, the experiment ended upon attainment of a learning criterion or after 256 trials in the absence of learning to criterion.

Figure 16
Mean Error and Number of Trials Completed for Each Participant

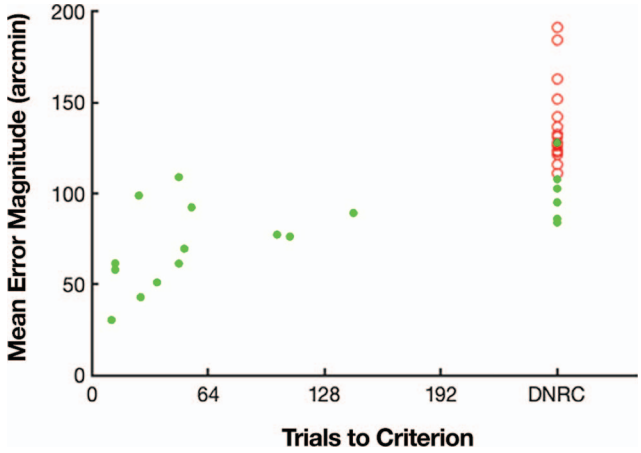

Note. Green, filled circles indicate participants in the constant curvature condition; red, open circles indicate participants in the constant acceleration condition. Did not reach the learning criterion (DNRC) indicates the number of trials (256) completed by participants who did not reach the learning criterion. See the online article for the color version of this figure.

little evidence of learning throughout training. Similarly, the direction of errors in the constant curvature condition showed no systematic bias, consistent with access to an underlying representation that allowed observers to locate the transition point between constant curvature segments. In contrast, the responses in the Euler spiral condition did not seem to converge on the transition point between two different spiral segments. In summary, this experiment furnishes evidence from a new paradigm that constant curvature, not simply any curvature regularity, is important for how the visual system represents contour shape.

The results of Experiment 2 also challenge some theories about the role of Euler spirals in visual shape perception. Kimia et al. (2003) proposed that the interpolating contour in shape completion has the form of an Euler spiral. However, participants appear to have no sensitivity to changes in the physical properties of two smoothly joined Euler spirals. We would expect better performance in the constant acceleration task if Euler spirals were a basic unit in form perception. On the other hand, participants' good performance in the constant curvature condition is consistent with a different theory about the shape of interpolating contours, namely that they are made up of two constant curvature segments (Ullman, 1976) or of a segment of zero curvature and a segment of constant curvature (Kellman & Shipley, 1991).

To our knowledge, no other work on contour or shape representation would make the prediction that motivated this experiment and which was confirmed by its results. This may not reflect an omission of any other work so much as a difference in the aspect of perceptual processing involved. Models of contour representation in which the point of highest curvature is most informative (e.g., Attneave, 1954; Feldman & Singh, 2005) do not predict that a change in curvature between two parts would be more salient than a change in curvature acceleration. In these models, the most surprising, and therefore salient, continuation of a contour is one in which the contour deviates most from the straight tangent direction (i.e., where curvature is highest). These models do not make an explicit prediction about how the visual system segments a contour, but the join point between two regions of constant curvature would not, in general, have higher average curvature than other points along the contour. In particular, these models would not predict the join point between two constant curvature segments to be most salient, since one segment will always have higher curvature than the other, and therefore any continuation along the higher curvature segment would be more surprising. These approaches would also not appear to furnish any basis for locating a change point between two constant curvature segments more reliably than locating a change point between two Euler spirals. Models that partition shapes based on curvature minima (e.g., Hoffman & Richards, 1983; Siddiqi et al., 1996; Singh et al., 1999) or curvature zero crossings (e.g., Mokhtarian & Mackworth, 1986) also cannot account for these data, as all tested contours were monotonic in Experiment 2.

As mentioned above, models that assess the informational value of points or that address the salience or segmentation of parts may act upon representations of contours already encoded. They may incorporate or serve ecological priorities that are subsequent to the initial representation of basic contour shape. The current results on contour shape representation remind us that the representation of the physical stimulus is not the same as the stimulus itself. Taking this a step further, the determination of physically or ecologically relevant parts of an object may involve factors that extend beyond basic representation of contour shape. Conversely, the "parts" in a constant curvature model of contour shape may function as an economical and effective structural description, but the function of such a representation is not to make explicit the ecologically or physically relevant parts of an object. Indeed, Wertheimer's (1923) classic introduction of the principle of good continuation involved demonstrations of continuous contours and how observers naturally break them into parts. Kellman et al. (2003) noted that part segmentation in these displays always involved a first-order or tangent discontinuity (a sharp corner), whereas changes in slope in a smooth contour did not disrupt perception of a single segment. All of the stimuli in both conditions in the present experiment involved contours that were "smooth" (having a continuous first derivative, when characterized as a function). In other words, the constant curvature hypothesis may provide a viable account of structural parts in contour encoding but it is not a theory of part decompositions that participants would spontaneously produce (e.g., De Winter & Wagemans, 2006).

This last observation coheres with the nature and results of Experiment 2 in another important way. Note that we did not predict that observers would naturally see break points in the contours presented at the start of the experiment. We performed a learnability study on the hypothesis that constant curvature constituents are not perceptually salient parts but that with some feedback, perceivers could come to access these parts that do exist in an underlying representation. The results showed little evidence of immediate perceptual salience but clear evidence that with a modest learning period and feedback, most observers readily became able to identify accurately the constant curvature segmentation points. This outcome was not observed in the control condition, despite a (formally) clear geometric criterion for segmentation of each curve into two Euler spirals.

Experiment 3

The preceding experiments furnished substantial new evidence supporting the idea of constant curvature encoding of contour shape. Experiment 3 tested another consequence of the hypothesis of constant curvature coding of contour shape, one more deeply embedded in the implications of the arclet model we introduced earlier.

As described, encoding of contour segments into regions of constant curvature, under the arclet theory, produces for each contour segment a description in terms of turning angle, scale, and extent. Two contours that differ only by a factor of scale have the same shape because the best fitting arclets for the two contours or contour segments will differ in scale but will match in the two parameters of turning angle and extent. Because a scale invariant shape code emerges naturally in this scheme (comprised of the match on the turning angle and extent parameters), we predicted that changes to the *extent* of a constant curvature segment would be more easily perceived than changes in scale that alter the absolute (mathematical) curvature of a segment. The reason is that differences in angular extent produced by arclet encoding unequivocally signal a shape difference between two contour segments, whereas a change in scale alone leads to encoding of two segments that match in terms of scale invariant shape coding. To more fully

motivate this prediction, we return to some details of the arclet model.

Constant curvature shape encoding in the arclet scheme makes a prediction about differential sensitivity to two types of differences between two contour segments. One way in which two constant curvature contour segments may differ is if they have the same shape but different scales. A scale difference does not change shape. Recall that the scale invariant code for a contour segment is {Θ, n}, where Θ describes the turning angle of the largest adequately fitting arclet, and *n* describes the number of arclet units or angular extent of the segment. Two segments differing only in scale would match on these two parameters. This match supports our natural perceptual tendency to see objects or contours as having the same shape, despite size differences. Indeed, in ordinary viewing, observers move closer or further away from objects. It is important to notice that such changes in viewing distance change the mathematical curvature at every point of a smooth contour; yet we do not see shape changing with variations in viewing distance. Likewise, simultaneously present objects of differing size but identical shapes are readily perceived to have a common shape. It is possible that these perceptual outcomes depend on recomputing of shape via adjusting mathematical curvature through rapid normalization and comparison processes, but the ease and obviousness of detecting shape apart from scale is highly consistent with the more direct encoding of scale invariant shape in the framework we have outlined.

One can run this argument in reverse, in a sense. Imagine two open contours, each made up of two constant curvature segments. In comparing these, if one of the segments in one of the contours is enlarged or reduced by some scale factor, this change might be relatively difficult to detect. This prediction would be counterintuitive on many other schemes for encoding contours, but in the proposed framework it arises from the fact that the scale of the best-fitting arclets is in fact extraneous to shape. A change in scale only for a contour or contour segment does not signal a change in the shape. On any scheme in which absolute curvature must first be detected, a contour segment can be compared at different scales only after normalization. In any such scheme, a scale change alone should provide a highly salient difference between stimuli.

For comparison, consider a different kind of change. If two constant curvature segments of contours being compared share the same scale but differ in the number of elements (extent), these would have differing shapes in our scheme, as they would match on the Θ and k parameters but not on *n*. Two such segments would not have the same scale invariant shape. Thus, a change in the *n* parameter might be highly detectable.

To summarize, changes in the angular extent between two contour segments should be readily detectable through a direct pairwise comparison of representations of contours. On the other hand, if one piece of a contour maintains its scale-invariant shape but is made smaller or larger, detection of this change may require recognition that the scale of the represented shape in one part of the contour has changed more or less than the scale at other parts. This more relational comparison should make shape changes of this kind more difficult to detect than changes to the angular extent of a part, which is directly comparable. We tested these predictions in Experiment 3.

Method

Participants

Initial results found $\eta_p^2 > .80$ for change type, suggesting a minimum of eight subjects for 80% power. Experiment 3 consisted of 12 naïve participants from UCLA who participated for course credit. Three additional psychophysical observers naive to the purpose of the experiment, and one author (PJK) also participated in the experiment and received no incentive for participation. All participants reported normal or corrected-to-normal vision. No participants were excluded. All procedures completed by participants in this study were approved by the UCLA Institutional Review Board.

Stimuli

Shapes presented to participants were composed of two smoothly connected constant curvature contour segments of opposite curvature polarity. Each shape was uniformly scaled to be no larger than 10.1° × 10.1°. Across stimuli, contour segment curvatures and lengths varied by up to 57%. For each shape, a comparison shape was generated. For half of the comparison shapes, the curvature of one of the segments was different, but the angular extents of both segments were the same. For the other half, the angular extent of one of the segments was different but the curvatures of both segments were the same (see Figure 17).

Although any change in scale results in an absolute curvature difference, the ratio of the curvatures within a shape is scale-invariant. Across the set of comparison stimuli, the magnitude of the change in this ratio was matched to the magnitude of the change in angular extent. Specifically, the magnitude of the changes was equated by matching the proportional differences between the two conditions. Proportional change magnitudes between 1.05:1 and 1.40:1 at increments of 0.05 were used.

Design and Procedure

Shape change type and magnitude were varied within subjects. The participants' task was to determine if the shapes of two figures were the same or different. On every trial two shapes were presented sequentially. The first shape was shown in the top left quadrant of the screen. The second shape was presented scaled (uniformly by 125% or 80%) and rotated (by 10°) relative to the other. The second shape was also moved to the bottom right quadrant of the screen. This position change, with the timing and spacing used, appeared sufficient to eliminate apparent motion. Participants were instructed to not consider overall scale, position, or orientation in their judgments of shape.

The beginning of each trial was signaled by a tone. After 250 ms, the first figure appeared and remained visible for 1,000 ms. Then, the figure was immediately replaced with its scaled, rotated counterpart which remained visible until a keypress was made to indicate whether the two figures' shapes were the same or different. After a keypress was made, feedback was given using two distinct tones for correct and incorrect responses and the next trial began. There were 64 practice trials and 640 experimental trials. Practice trials exposed participants to a representative sample of the different trial conditions of the experiment. A schematic of Experiment 3 trials is shown in Figure 18.

Figure 17
Schematic for the Two Kinds of Stimulus Changes in Experiment 3

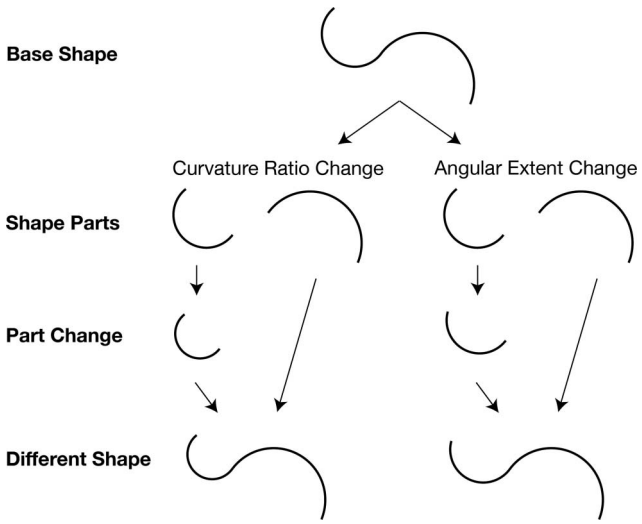

Note. Left: Changes in curvature ratio. Right: Changes to the angular extent of one of the constant curvature segments.

Results

Sensitivity (d') for all difference magnitudes for the two change types is shown in Figure 19. A hit was defined as a correct identification of a shape change, and a false alarm was defined as a reported shape change when the first and second shapes were the same. False alarm rates (32%) were therefore the same for both curvature ratio and angular extent changes. Sensitivity to changes increased with change magnitude and was markedly higher, across change magnitudes, for changes in the angular extent of the contour relative to changes in curvature ratio. These patterns were confirmed by the analyses. Sensitivity was compared in a 2 (angular extent change vs. curvature ratio change) $\times$ 8 (shape difference magnitude) repeated measures ANOVA, with both factors tested within subjects. There was a main effect for shape difference magnitude, $F(7, 15) = 35.57$, $p < .01$, adjusted $\eta_p^2 = .56$. Sensitivity was higher for changes in a contour's angular extent than for changes in a contour's curvature ratio, $F(1, 15) = 133.72$, $p < .01$, adjusted $\eta_p^2 = .89$. There was also a significant interaction between shape difference magnitude and shape change type, indicating that the difference in sensitivity between angular extent and curvature ratio was larger when shape difference magnitudes were also large, $F(7, 15) = 10.30$, $p < .01$, adjusted $\eta_p^2 = .13$.

Discussion

The results of Experiment 3 are consistent with predictions of the arclet theory that contours are represented as joined regions of constant curvature. In the representation of a contour with more than one constant curvature segment, each region would have a position, orientation, angular extent, curvature, and size. Turning angles and angular extents of the largest well-fitting arclets provide a scale-invariant code for the shape of a contour segment, irrespective of size. When the angular extent of one segment is changed, the scale-invariant code changes. When the size of one segment is changed, the scale-invariant codes for each segment remain the same. Considering the whole contour, made of two segments, what has changed is the ratio of their scales in the scale-specific code.

On the basis of these theoretical ideas, we predicted that a change in angular extent would be more easily detectable than a change in the size of a segment. The task given was to compare shapes between the initial stimulus and a scaled, rotated display. If

Figure 18
Stimuli and Procedure for Experiment 3

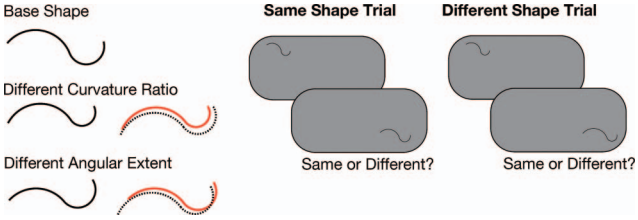

Note. A base shape (left, top) was generated by smoothly joining two constant curvature segments with opposite curvature polarity. Shape comparisons were made to either the same shape, a shape in which the curvature value of one of the segments changed, but the angular extent remained the same (left, middle), or a shape in which the curvature values remained the same, but the angular extent of one of the segments changed (left, bottom). Different shapes are shown superimposed on the base shape, with the base shape in dotted lines and the compared shape in red. A sample same and different trial are shown (right). Note that scale varied across presentations to eliminate change in size as a cue for the same/different task. See the online article for the color version of this figure.

Figure 19
Results of Experiment 3

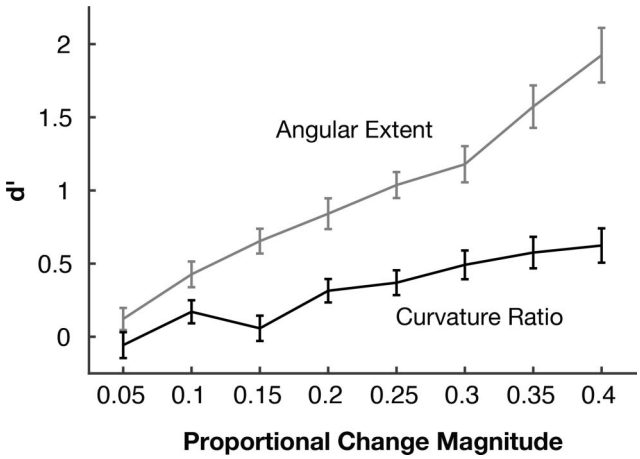

Note. Mean sensitivity to changes in shape are shown as a function of the magnitude of the change. Proportional change magnitude is the size of the change between two shape presentations on a different shape trial, expressed as a proportion of the magnitude of the value that changed (angular extent or curvature ratio) from presentation 1 to presentation 2. Error bars indicate ± 1 standard error of the mean. Angular Extent changes are shown in the top, gray line. Curvature ratio changes are shown in the bottom, black line.

each two-segment contour was encoded in our constant curvature scheme, changes to a segment's angular extent would have been detectable directly, from comparison of the corresponding segments' scale-invariant codes. In contrast, a change from the initial stimulus that involved a differential scaling of one segment relative to the other would have been detectable only by first calculating the ratio of scale-specific curvatures within a shape, and then comparing these ratios across shapes. We predicted that this more complicated comparison would result in lower change-detection sensitivity. Importantly, this is not a prediction that would be made by representational schemes that do not posit a decomposition into parts of constant curvature. For example, in Figure 17, the difference in lengths between the two parts is larger in the change to curvature ratio condition than in the change to angular extent condition.

General Discussion

The primary purpose of this work was to assess the idea that representation of 2D contour shape involves constant curvature segments as primitives. We hypothesized that if the visual system uses constant curvature segments in their representations, it should have special capabilities to extract, encode, or compare contours with constant curvature over and above its capabilities for other contours. To ensure the generality of this hypothesis, we tested it using three very different methods.

In Experiment 1, we used a path detection paradigm to test whether paths with constant turning angle were easier to detect among oriented distractors than paths with varying turning angles. Detection performance for the constant turning angle paths was reliably higher than detection performance for paths with alternating turning angle (Experiment 1a) and paths with a consistent direction of turning but varying turning angle (Experiment 1b). The superior performance for constant turning angle paths did not depend on the magnitude of the curvature: Performance was higher for CC paths whether they had higher or lower average curvature than the mixed curvature paths. Participants' better performance for constant curvature paths is consistent with the presence of visual mechanisms sensitive to constant curvature. Because both constant curvature and nonconstant curvature paths had the same relations between pairs of oriented line segments, the results indicate that detectors with constant turning angles along longer portions of a potential contour facilitate detection. This would naturally follow from cooperative facilitation of adjacent arclets with the same turning angle and curvature polarity, leading to extraction of constant curvature representations of contours.

In Experiment 2, we compared participants' ability to learn to divide a contour into two constant curvature parts with their ability to learn to divide a contour into two Euler spiral parts (i.e., two segments with constantly accelerating curvature). We hypothesized that participants would be able to learn to segment the constant curvature fragments without much difficulty, but that the constant acceleration fragments would be difficult to segment because the visual system does not encode contour shape with higher-order curvature relations. Twelve out of 18 participants tested with constant curvature fragments learned to accurately segment the contour well enough to meet a demanding criterion of 10 consecutive precise segmentations, and those who met criterion did so much earlier than the designated cutoff point of 256 trials. By contrast, no subject tested with Euler spiral fragments learned to segment to criterion within that timeframe. These findings suggest that contours are automatically encoded as a set of constant curvature primitives. These primitives are descriptors for the shapes of continuous contours rather than markers for segmentation of visual input into separate objects or separate contours. These results, showing that a brief learning opportunity with feedback allows observers to access components in a contour representation, are consistent with the existence of constant curvature coding of contours.

Experiment 3 tested subjects' sensitivity to changes in a contour made of two parts. We found robustly greater sensitivity to a change in the angular extent of a part than to a change in the mathematical curvature of a part. We described a particular approach to contour curvature encoding, the arclet framework, and described how this framework, when implemented at multiple scales, obtains both a scale-invariant and scale-specific representation of the shapes. In the arclet framework, each part is represented by a curvature (turning angle), an angular extent, and a scale factor of the largest scale arclets that provide an adequate fit. Thus, for two contour components of the same shape but differing size, the two parameters that define shape—turning angle and extent—remain invariant. Based on this property, we predicted that when sequentially viewing and comparing two contours, a change consisting of the scaling of one component of the contour would be a relatively difficult change to detect. In our paradigm, the comparison stimulus was always globally scaled and rotated relative to the initial stimulus. Within the arclet framework, detecting that the relative size of two components in the comparison stimulus differed from the initial stimulus would require computation relating to scale-specific representations and their relations between contour components. In contrast, a change in angular extent from target to comparison stimulus should be directly detectable as a change in the scale invariant representation. These predictions followed from the arclet framework but would not be expected in the absence of encoding of contour shape in terms of constant curvature segments. The results of Experiment 3 clearly and robustly supported these predictions: Participants consistently showed higher sensitivity to changes in a segment's angular extent than for equivalent magnitude changes to segment curvature.

Taken together, the findings of these experiments provide strong and converging evidence for the idea that segments of constant curvature are the building blocks of contour shape representations. The evidence shows the relevance of constant curvature contour representations in detection of stimuli in noise, in the ease of learning to segment contours into parts, and in the detection of changes in contour shapes. These results all indicate that the visual system performs more accurately perceptual tasks that would be predicted to be facilitated by an underlying model in which constant curvature segments are privileged in the encoding of shape. In each of the experiments, many local properties, such as the total amount of curvature, the sign of curvature, the degree of curvature change, or overall change were equated across stimulus conditions.

Experimental Results and the Arclet Framework

The present results converge in furnishing evidence that contour shape representation utilizes constant curvature segments. In addition to this general idea, we proposed the arclet model to explain

how constant curvature representations are obtained. This framework indicates how curvature may be extracted from known, noncurved detectors; how the visual system gets from neural units coding contrast to symbolic representations of shape; and how notions of scale invariance, such as the idea that the same shape is readily recognized for objects of differing sizes, may be natural consequences of this framework.

We have sketched the basic framework for how curvature detectors (arclets), built from the outputs of oriented units, could allow extraction of constant curvature representations from arbitrary contour shapes. Working computational models for this process have been developed (Garrigan & Kellman, 2011), but further work is needed to build a fully working model from plausible neural-style inputs. Such a model would comprise an existence proof of visual processes that begin with subsymbolic inputs (oriented contrast detectors) and produce symbolic representations of shape (connected contour tokens with shape descriptions in a scheme that readily allows comparisons across scale, orientation, etc.). A particularly crucial step is to specify how arclet responses at different scales are integrated, and specifically how to identify the largest scale that suitably fits a contour of a given curvature and size.

Although the present tests should not be considered exhaustive by any means, the experiments here are highly consistent with the predictions of the arclet approach. As noted, the results of Experiment 3 tested a specific prediction of the model that would not be expected on other grounds. The results of Experiment 1 are also consistent with the activation of higher-order neural detectors that are tuned to constant curvature segments.

Relations to Other Approaches and Phenomena

It is worth noting that none of the primary predictions in any of the experiments would be straightforward consequences of other approaches to object and shape perception. Superior performance with constant curvature contours is unlikely to be explainable by the visual system's greater familiarity with constant curvature stimuli. Scene analysis has found that circular contours do not occur in natural environments more frequently than other smooth, closed contours (Chow et al., 2002). Rather, we suggest that the visual system encodes all contours in terms of constant curvature primitives. Contours that already have constant curvature can be encoded with relatively few elements and are therefore simpler (Garrigan & Kellman, 2011). More commonly, contours with nonconstant curvature must be segmented into regions of near-constant curvature and encoded with more elements, resulting in a more complex representation that might give rise to lower performance on perceptual tasks, particularly when exposure time is limited or the perceptual judgment is difficult.

The constant curvature theory of contour representation predicts that contour regions are represented by different primitives than part segmentation theories, such as codon theory (Richards & Hoffman, 1985), limbs and necks (Siddiqi et al., 1996), the shortcut rule (Singh et al., 1999), or an integrated account of all three (De Winter & Wagemans, 2006). In these theories, parts are segmented based on their status between contour minima or inflection points. In our experiments, a difference in representational complexity was observed between contour regions that had neither minima nor curvature sign changes.

We do not see these theories as being in opposition to the constant curvature theory; rather, theories of part segmentation may not capture the most basic primitives of contour representation. The contour segmentations produced by these theories have strong predictive power for certain perceptual tasks, such as unguided segmentation of a contour into parts (De Winter & Wagemans, 2006). These tasks aim at identifying the perceptually salient parts of objects, often corresponding to semantic properties like a limb, handle, or neck. The model of constant curvature primitives we propose does not aim to identify perceptually salient parts in this sense but to provide an economical representation of contour shape in the first place. Ecologically, the purposes of seeing salient parts, such as grasping, taking things apart, predicting where things will break, inferring functional properties, and so forth, are not the same as encoding a representation of the contour itself. The former may be more cognitive in nature, as suggested by the finding that segmentations differed if an object was familiar or not (De Winter & Wagemans, 2006), while the latter is entirely perceptual and deals with the first transition from a literal representation of the contour as might exist in a visual icon to an abstract representation of shape. Our results show that with modest amounts of learning, aspects of this representation can become perceptually noticeable.

One notion of shape representation that has some conceptual overlap with the constant curvature model we propose is structural information theory (SIT). According to SIT, the visual system aims to find the simplest symbolic code for a contour based on rules of iteration, alternation, and symmetry (van der Helm, 2011; van der Helm et al., 1992; Wagemans, 2015). The symbolic representation of a contour region as a single-curvature segment in our theory is similar to a certain kind of iterative code in SIT. We describe a constant curvature segment by a turning angle, scale, and extent, which could easily be formalized in an SIT framework as $N \times$ [s, k], where s and k are the scale and turning angle, and N is the extent specifying the number of iterations of a certain scale and turning angle in a given curvature segment.

The constant curvature model and SIT appear to operate on different levels and address different problems, however. Specifically, the constant curvature model encompasses the transition from subsymbolic to symbolic representations of shape. SIT is a very general theory, attempting to account for a wide range of regularities in visual arrays (van der Helm et al., 1992). In SIT, the contour is taken as a given, already represented by, and partitioned into, a symbol series. The goal of SIT is to find efficient ways to code the symbol series, whereas in the constant curvature model, the goal is to represent something with initially no symbolic structure. The constant curvature model also predicts that one kind of coding rule (iteration) is simpler to encode than the others. This was borne out in Experiment 1a, where a contour path of constant turning angle was more efficiently encoded than a contour path of alternating turning angle, both of which would have simple symbolic codes in SIT. Finally, the constant curvature model specifies a way for things that are not actually redundant in the stimulus to be represented compactly. In SIT, a simple symbolic code could be formed for a contour region that does have constant curvature, but not for a contour region with similar but varied curvatures. Ecological evidence points to contours in visual scenes not being truly cocircular (Chow et al., 2002). The constant curvature model offers a way for a contour region made up of different but similar

curvatures to be organized together into a region of singular curvature. In this article, we have focused on people's ability to encode contours that have constant curvature, but the theory in principle extends to more naturalistic contours which do not have extended regions of constant curvature, and there is experimental evidence supporting the constant curvature encoding of such contours (Baker, 2020; Garrigan & Kellman, 2011).

Apart from comparisons to other general approaches, each of the experiments presented here contrasted the predictions of a constant curvature approach to what might be expected on the grounds of a plausible alternative model. In Experiment 1, the benefits of constant curvature exceeded what would be predicted merely by constant polarity of orientation change along the path. The advantages predicted for constant curvature paths also went beyond known characteristics of relations previously hypothesized to be relevant to the "association field" framework, or similarly, contour relatability in contour interpolation. In Experiment 2, the prediction of the constant curvature model not only fared better than the prediction of accelerating spiral models (e.g., Kimia et al., 2003), but segmentation of a contour into two segments based on the latter seemed, within the context of our experiment, unlearnable. Finally, the prediction in Experiment 3, and the evidence of superior detection of change for angular extent of a contour segment versus relative size (when scale invariant curvature was preserved) would be hard to understand from the perspective of any of many models in computational or biological vision that rely on mathematical curvature and normalization based on overall contour or figure size.

Limitations of the Present Work

A limitation of the present work is that the constant curvature hypothesis models only a description of 2D contour shape. This is an important aspect of visual perception, and no doubt understanding shape perception requires multiple levels of description, but it would be useful to investigate the relationship between 2D contour shape and 3D object perception, a subject which has already received some research efforts (Koenderink, 1984; Li et al., 2009; Qian et al., 2019). It is also likely that even closed 2D shapes require a more sophisticated description of how the regions between segments are represented than has been considered here. On the other hand, it also seems quite likely that higher-level descriptions of 2D and 3D objects, such as symmetry, require as inputs initial symbolic descriptions of projected (2D) contour shape, as addressed here.

Another exciting possibility is that the core notion of constant curvature could also be extended to efficient description of 3D surface shape. It might be possible, for example, to define surface shape at discrete patches according to the Koenderink Shape Index (Koenderink, 1989), and then recode adjacent patches of surface in a manner analogous to capturing approximate segments of constant contour curvature as in our computational model (Garrigan & Kellman, 2011). Alternatively, the constant curvature 2D shape description could be an input to other kinds of volumetric shape models (e.g., Biederman, 1987; Marr & Nishihara, 1978).

It is hoped that the present work advances our understanding of shape perception, a fundamental but poorly understood area in the study of visual perception. The results also open up many new questions and potential lines of investigation. In a larger perspective, these results, and the theoretical framework described, may also provide an example of how theories of middle and high-level vision may bridge the gap between subsymbolic encoding of, for example, contrast and retinal orientation, and symbolic descriptions of an object's shape. Taken together with ecological evidence for the usefulness of constant curvature primitives for approximating closed, smooth contours and neurophysiological evidence of sensitivity to constant curvature, the current findings implicate constant curvature segments as fundamental primitives of abstract shape representation.

Context

How we perceive and represent shape is among the most fundamental problems of perception and cognition. Core problems of shape perception were the central focus of classic work by Gestalt psychologists a century ago, but attaining modern scientific explanations in terms of computational processes and neural coding have remained elusive. Much of our prior work has focused on understanding visual perception of contours, surfaces, and objects; our work and that of others have often implicated the relational, abstract nature of perception. It also partakes of a broadly important scientific mystery. Considerable progress has been made in understanding neural units in early cortical processing that respond to local, oriented contrast. Meanwhile, work in higher level vision has revealed much about perceptual processing of contours, objects, and shape, usually starting from abstract, mathematical descriptions of stimuli. Yet little or no work bridges the gap, indicating how from initial, local, transient, subsymbolic responses to contrast we perceive shape and achieve economical, versatile, abstract shape representations. The findings and theory presented here suggest how this transition may work for contour shape; the transition depends on oriented units but uses linkages among them to code contour shape into constant curvature segments. These segments, joined together, describe contours. Besides offering an account of contour shape, an important aspect of shape perception, the work puts forward a means of bridging from early, transient, subsymbolic activations to more enduring and abstract representations that may help suggest how this bridge is built in other areas of visual perception.

References

Attneave, F. (1954). Some informational aspects of visual perception. *Psychological Review*, *61*(3), 183–193. https://doi.org/10.1037/h0054663

Ayzenberg, V., & Lourenco, S. F. (2019). Skeletal descriptions of shape provide unique perceptual information for object recognition. *Scientific Reports*, *9*(1), 1–13. https://doi.org/10.1038/s41598-019-45268-y

Baker, N. (2020). *Curve-based shape representation in visual perception*. University of California.

Baker, N., & Kellman, P. J. (2018). Abstract shape representation in human visual perception. *Journal of Experimental Psychology: General*, *147*(9), 1295–1308. https://doi.org/10.1037/xge0000409

Baker, N., Lu, H., Erlikhman, G., & Kellman, P. J. (2018). Deep convolutional networks do not classify based on global object shape. *PLoS Computational Biology*, *14*(12), e1006613. https://doi.org/10.1371/journal.pcbi.1006613

Barenholtz, E., Cohen, E. H., Feldman, J., & Singh, M. (2003). Detection of change in shape: An advantage for concavities. *Cognition*, *89*(1), 1–9. https://doi.org/10.1016/S0010-0277(03)00068-4

Barsalou, L. W., & Wiemer-Hastings, K. (2005). Situating abstract concepts. In D. Pecher & R. Zwaan (Eds.), *Grounding cogmtwn: The role of perception and action in memory, language, and thought* (pp. 129–163). Cambridge University Press. https://doi.org/10.1017/CBO9780511499968.007

Bell, J., Badcock, D. R., Wilson, H., & Wilkinson, F. (2007). Detection of shape in radial frequency contours: Independence of local and global form information. *Vision Research*, *47*(11), 1518–1522. https://doi.org/10.1016/j.visres.2007.01.006

Bertamini, M. (2001). The importance of being convex: An advantage for convexity when judging position. *Perception*, *30*(11), 1295–1310. https://doi.org/10.1068/p3197

Biederman, I. (1987). Recognition-by-components: A theory of human image understanding. *Psychological Review*, *94*(2), 115–147. https://doi.org/10.1037/0033-295X.94.2.115

Carrigan, S. B. (2020). *Promiscuity and pruning: Investigations of a two-stage theory of contour interpolation* [Doctoral dissertation]. Retrieved from https://escholarship.org/uc/item/41j2x52p

Carrigan, S. B., & Kellman, P. J. (2020). Combined effects of multiple scene cues on the perceptual strengths of promiscuously interpolated contours. *Journal of Vision*, *20*(11), 1276. https://doi.org/10.1167/jov.20.11.1276

Chow, C. C., Jin, D. Z., & Treves, A. T. (2002). Is the world full of circles? *Journal of Vision*, *2*, 571–576. https://doi.org/10.1167/2.8.4

Collin, C. A., Liu, C. H., Troje, N. F., McMullen, P. A., & Chaudhuri, A. (2004). Face recognition is affected by similarity in spatial frequency range to a greater degree than within-category object recognition. *Journal of Experimental Psychology: Human Perception and Performance*, *30*(5), 975–987. https://doi.org/10.1037/0096-1523.30.5.975

Cumming, G. (2012). *Understanding the new statistics: Effect sizes, confidence intervals, and meta-analysis*. Routledge.

Davis, L. S. (1977). Understanding shape: Angles and sides. *IEEE Transactions on Computers*, *C-26*(3), 236–242.

De Valois, R. L., & De Valois, K. K. (1980). Spatial vision. *Annual Review of Psychology*, *31*(1), 309–341.

De Winter, J., & Wagemans, J. (2006). Segmentation of object outlines into parts: A large-scale integrative study. *Cognition*, *99*(3), 275–325.

Dubinksiy, A., & Zhu, S. C. (2003). A multi-scale generative model for animate shapes and parts. *Proceedings of the Ninth IEEE International Conference on Computer Vision*, *1*, 249–256. https://doi.org/10.1109/ICCV.2003.1238350

Elder, J. H., Oleskiw, T. D., Yakubovich, A., & Peyré, G. (2013). On growth and formlets: Sparse multi-scale coding of planar shape. *Image and Vision Computing*, *31*(1), 1–13.

Farah, M. J., Rochlin, R., & Klein, K. L. (1994). Orientation invariance and geometric primitives in shape recognition. *Cognitive Science*, *18*, 325–344.

Feldman, J., & Singh, M. (2005). Information along contours and object boundaries. *Psychological Review*, *112*(1), 243.

Feldman, J., & Singh, M. (2006). Bayesian estimation of the shape skeleton. *Proceedings of the National Academy of Sciences of the United States of America*, *103*(47), 18014–18019.

Field, D. J., Hayes, A., & Hess, R. F. (1993). Contour integration by the human visual system: Evidence for a local "association field." *Vision Research*, *33*(2), 173–193.

Garrigan, P., & Kellman, P. J. (2011). The role of constant curvature in 2-D contour shape representations. *Perception*, *40*(11), 1290–1308. https://doi.org/10.1068/p6970

Gibson, J. J. (1966). *The senses considered as perceptual systems*. Houghton Mifflin.

Gibson, J. J. (1979). *The ecological approach to visual perception*. Psychology Press.

Graham, S. A., Kilbreath, C. S., & Welder, A. N. (2004). Thirteen-month-olds rely on shared labels and shape similarity for inductive inferences. *Child Development*, *75*(2), 409–427. https://doi.org/10.1111/j.1467-8624.2004.00683.x

Hess, R., & Field, D. (1999). Integration of contours: New insights. *Trends in Cognitive Sciences*, *3*(12), 480–486. https://doi.org/10.1016/S1364-6613(99)01410-2

Hess, R., Hayes, A., & Field, D. (2000). The roles of polarity and symmetry in the perceptual grouping of contour fragments. *Spatial Vision*, *13*(1), 51–66. https://doi.org/10.1163/156856800741018

Hoffman, D. D., & Richards, W. (1983). Parts of recognition. *Cognition*, *18*(1–3), 65–96.

Hubel, D. H., & Wiesel, T. N. (1968). Receptive fields and functional architecture of monkey striate cortex. *The Journal of Physiology*, *195*(1), 215–243. https://doi.org/10.1113/jphysiol.1968.sp008455

Kass, M., Witkin, A., & Terzopoulos, D. (1988). Snakes: Active contour models. *International Journal of Computer Vision*, *1*(4), 321–331. https://doi.org/10.1007/BF00133570

Kellman, P. J., Erlikhman, G., & Carrigan, S. B. (2016). Is there a common mechanism for path integration and illusory contour formation? *Journal of Vision*, *16*(12), 311. https://doi.org/10.1167/16.12.311

Kellman, P. J., & Garrigan, P. (2007). Segmentation, grouping, and shape: Some Hochbergian questions. In M. A. Peterson & B. Gillam (Eds.), *Sedgwick, perception: Essays in honor of Julian Hochberg* (pp. 542–554). Oxford University Press.

Kellman, P. J., Garrigan, P., & Erlikhman, G. (2013). Challenges in understanding visual shape perception and representation: Bridging subsymbolic and symbolic coding. In S. Dickinson & Z. Pizlo (Eds.), *Shape perception in human and computer vision* (pp. 249–274). Springer. https://doi.org/10.1007/978-1-4471-5195-1_18

Kellman, P. J., Garrigan, P., Kalar, D., & Shipley, T. F. (2003). Good continuation and relatability: Related but distinct principles. *Journal of Vision*, *3*(9), 120a. https://doi.org/10.1167/3.9.120

Kellman, P. J., Garrigan, P., & Shipley, T. F. (2005). Object interpolation in three dimensions. *Psychological Review*, *112*(3), 586–609. https://doi.org/10.1037/0033-295X.112.3.586

Kellman, P. J., Guttman, S. E., & Wickens, T. D. (2001). Geometric and neural models of object perception. In T. F. Shipley & P. J. Kellman (Eds.), *Advances in psychology, 130. From fragments to objects: Segmentation and grouping in vision* (pp. 183–245). Elsevier Science. https://doi.org/10.1016/S0166-4115(01)80027-3

Kellman, P. J., & Shipley, T. F. (1991). A theory of visual interpolation in object perception. *Cognitive Psychology*, *23*(2), 141–221. https://doi.org/10.1016/0010-0285(91)90009-D

Kimia, B. B., Frankel, I., & Popescu, A. M. (2003). Euler spiral for shape completion. *International Journal of Computer Vision*, *54*(1–3), 159–182.

Koenderink, J. J. (1984). What does the occluding contour tell us about solid shape? *Perception*, *13*, 321–330. https://doi.org/10.1068/p130321

Koenderink, J. J. (1989). *Solid shape*. MIT Press.

Koffka, K. (1935). *Principles of Gestalt psychology*. Routledge.

Kohler, W. (1929). *Gestalt psychology*. Horace Liverwright.

Kovacs, I., & Julesz, B. (1993). A closed curve is much more than an incomplete one: Effect of closure in figure-ground segmentation. *Proceedings of the National Academy of Sciences of the United States of America*, *90*(16), 7495–7497. https://doi.org/10.1073/pnas.90.16.7495

Kulikowski, J. J., & Bishop, P. O. (1981). Linear analysis of the responses of simple cells in the cat visual cortex. *Experimental Brain Research*, *44*(4), 386–400. https://doi.org/10.1007/BF00238831

Lakens, D. (2013). Calculating and reporting effect sizes to facilitate cumulative science: A practical primer for t-tests and ANOVAs.

Frontiers in Psychology, *4*, 863. https://doi.org/10.3389/fpsyg.2013.00863

Lederman, S. J., & Wing, A. M. (2003). Perceptual judgement, grasp point selection and object symmetry. *Experimental Brain Research*, *152*(2), 156–165. https://doi.org/10.1007/s00221-003-1522-5

Li, Y., Pizlo, Z., & Steinman, R. M. (2009). A computational model that recovers the 3D shape of an object from a single 2D retinal representation. *Vision Research*, *49*(9), 979–991. https://doi.org/10.1016/j.visres.2008.05.013

Lowet, A. S., Firestone, C., & Scholl, B. J. (2018). Seeing structure: Shape skeletons modulate perceived similarity. *Attention, Perception & Psychophysics*, *80*(5), 1278–1289. https://doi.org/10.3758/s13414-017-1457-8

Marĉelja, S. (1980). Mathematical description of the responses of simple cortical cells. *JOSA*, *70*(11), 1297–1300. https://doi.org/10.1364/JOSA.70.001297

Marcus, G. F. (2001). *The algebraic mind: Reflections on connectionism and cognitive science*. MIT Press. https://doi.org/10.7551/mitpress/1187.001.0001

Marr, D. (1982). *Vision: A computational investigation into the human representation and processing of visual information*. Freeman.

Marr, D., & Nishihara, H. K. (1978). Representation and recognition of the spatial organization of three-dimensional shapes. *Proceedings of the Royal Society of London, Series B: Biological Sciences*, *200*, 269–294. https://doi.org/10.1098/rspb.1978.0020

Mokhtarian, F., & Mackworth, A. (1986). Scale-based description and recognition of planar curves and two-dimensional shapes. *IEEE Transactions on Pattern Analysis and Machine Intelligence, PAMI*, *8*(1), 34–43. https://doi.org/10.1109/TPAMI.1986.4767750

Mordkoff, J. T. (2019). A simple method for removing bias from a popular measure of standardized effect size: Adjusted partial eta squared. *Advances in Methods and Practices in Psychological Science*, *2*(3), 228–232. https://doi.org/10.1177/2515245919855053

Murphy, G. L., & Brownell, H. H. (1985). Category differentiation in object recognition: Typicality constraints on the basic category advantage. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *11*(1), 70–84. https://doi.org/10.1037/0278-7393.11.1.70

Neri, P. (2018). The empirical characteristics of human pattern vision defy theoretically-driven expectations. *PLoS Computational Biology*, *14*(12), e1006585. https://doi.org/10.1371/journal.pcbi.1006585

Overlan, M. C., Jacobs, R. A., & Piantadosi, S. T. (2017). Learning abstract visual concepts via probabilistic program induction in a language of thought. *Cognition*, *168*, 320–334. https://doi.org/10.1016/j.cognition.2017.07.005

Parent, P., & Zucker, S. W. (1989). Trace inference, curvature consistency, and curve detection. *IEE Transactions on Pattern Analysis and Machine Intelligence*, *11*(8), 823–839. https://doi.org/10.1109/34.31445

Pasupathy, A., & Connor, C. E. (2001). Shape representation in area V4: Position-specific tuning for boundary conformation. *Journal of Neurophysiology*, *86*(5), 2505–2519. https://doi.org/10.1152/jn.2001.86.5.2505

Pasupathy, A., & Connor, C. E. (2002). Population coding of shape in area V4. *Nature Neuroscience*, *5*(12), 1332–1338. https://doi.org/10.1038/972

Pettet, M. W. (1999). Shape and contour detection. *Vision Research*, *39*, 551–557. https://doi.org/10.1016/S0042-6989(98)00130-8

Pomerantz, J. R., & Portillo, M. C. (2011). Grouping and emergent features in vision: Toward a theory of basic Gestalts. *Journal of Experimental Psychology: Human Perception and Performance*, *37*(5), 1331–1349. https://doi.org/10.1037/a0024330

Pomerantz, J. R., Sager, L. C., & Stoever, R. J. (1977). Perception of wholes and of their component parts: Some configural superiority effects. *Journal of Experimental Psychology: Human Perception and Performance*, *3*(3), 422–435. https://doi.org/10.1037/0096-1523.3.3.422

Qian, Y., Ramalingam, S., & Elder, J. H. (2019). LS3D: Single-view gestalt 3D surface reconstruction from Manhattan line segments. In C. Jawahar, H. Li, G. Mori, & K. Schindler (Eds.), *Computer Vision – ACCV 2018. ACCV 2018. Lecture Notes in Computer Science, 11364* (pp. 399–416). Springer. https://doi.org/10.1007/978-3-030-20870-7_25

Rezanejad, M., & Siddiqi, K. (2013). Flux graphs for 2D shape analysis. In S. Dickinson & Z. Pizlo (Eds.), *Shape perception in human and computer vision* (pp. 41–54). Springer. https://doi.org/10.1007/978-1-4471-5195-1_3

Richards, W., & Hoffman, D. D. (1985). Codon constraints on closed 2D shapes. *Computer Vision Graphics and Image Processing*, *31*(3), 265–281. https://doi.org/10.1016/0734-189X(85)90031-3

Ringach, D. L. (2002). Spatial structure and symmetry of simple-cell receptive fields in macaque primary visual cortex. *Journal of neurophysiology*, *88*(1), 455–463.

Rossion, B., & Pourtois, G. (2004). Revisiting Snodgrass and Vanderwart's object pictorial set: The role of surface detail in basic-level object recognition. *Perception*, *33*(2), 217–236. https://doi.org/10.1068/p5117

Schmidtmann, G., Jennings, B. J., & Kingdom, F. A. (2015). Shape recognition: Convexities, concavities and things in between. *Scientific Reports*, *5*(1), 1–11. https://doi.org/10.1038/srep17142

Siddiqi, K., Tresness, K. J., & Kimia, B. B. (1996). Parts of visual form: Psychophysical aspects. *Perception*, *25*(4), 399–424. https://doi.org/10.1068/p250399

Sigman, M., Cecchi, G. A., Gilbert, C. D., & Magnasco, M. O. (2001). On a common circle: Natural scenes and Gestalt rules. *Proceedings of the National Academy of Sciences of the United States of America*, *98*, 1935–1940. https://doi.org/10.1073/pnas.98.4.1935

Simon, H. A., & Newell, A. (1976). Computer science as empirical inquiry: Symbols and search. *Communications of the ACM*, *19*(3), 11–126.

Singh, M., Seyranian, G. D., & Hoffman, D. D. (1999). Parsing silhouettes: The short-cut rule. *Perception & Psychophysics*, *61*(4), 636–660. https://doi.org/10.3758/BF03205536

Sperling, G. (1960). The information available in brief visual presentations. *Psychological Monographs*, *74*(11), 1–29. https://doi.org/10.1037/h0093759

Tarr, M. J., & Bulthoff, H. H. (1995). Is human object recognition better described by geon structural descriptions or by multiple views? Comment on Biederman and Gerhardstein (1993). *Journal of Experimental Psychology: Human Perception and Performance*, *21*(6), 1494–1505. https://doi.org/10.1037/0096-1523.21.6.1494

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

Ullman, S. (2007). Object recognition and segmentation by a fragment-based hierarchy. *Trends in Cognitive Sciences*, *11*(2), 58–64. https://doi.org/10.1016/j.tics.2006.11.009

van der Helm, P. A. (2011). Bayesian confusions surrounding simplicity and likelihood in perceptual organization. *Acta Psychologica*, *138*(3), 337–346. https://doi.org/10.1016/j.actpsy.2011.09.007

van der Helm, P. A., Van Lier, R. J., & Leeuwenberg, E. J. (1992). Serial pattern complexity: Irregularity and hierarchy. *Perception*, *21*(4), 517–544. https://doi.org/10.1068/p210517

Wagemans, J. (2015). Structural information theory: The simplicity of visual form, simplicity in vision: A multidisciplinary account of perceptual organization, psychology of touch and blindness, psychology of touch and blindness. *Perception*, *44*, 222–227. https://doi.org/10.1068/p4402rvw

Welder, A. N., & Graham, S. A. (2001). The influence of shape similarity and shared labels on infants' inductive inferences about nonobvious

object properties. *Child Development*, *72*(6), 1653–1673. https://doi.org/10.1111/1467-8624.00371

Wertheimer, M. (1923). Laws of organization in perceptual forms. In W. D. Ellis (Ed.), *A source book of Gestalt psychology* (pp. 71–88). Kegan Paul, Trench, Trubner & Company. https://doi.org/10.1037/11496-005

Zahn, C. T., & Roskies, R. Z. (1972). Fourier descriptors for plane closed curves. *IEEE Transactions on Computers*, *100*(3), 269–281. https://doi.org/10.1109/TC.1972.5008949

Zhao, Y., & Belkasim, S. (2012). Multiresolution Fourier descriptors for multiresolution shape analysis. *IEEE Signal Processing Letters*, *19*(10), 692–695. https://doi.org/10.1109/LSP.2012.2210040

Received April 6, 2020
Revision received October 6, 2020
Accepted October 7, 2020 ■