Chapter 5
Perceptual Organization of Shape

James H. Elder

5.1 Introduction

Computing the correct bounding contours of objects in complex natural scenes is generally thought to be one of the harder computer vision problems, and the state of the art is still quite far from human performance, when human subjects are given an arbitrary amount of time to delineate the shape of these contours [1]. How do we explain this gap? One possibility is that, when given enough time, humans fall back on high-level, deliberative reasoning processes. If this is true, then it is possible that when faced with detecting and recognizing objects in an active vision timeframe (hundreds of milliseconds), we may rely upon a simpler "bag of tricks", using appearance cues such as texture and color to discriminate objects.

To address this possibility, let us consider the specific task of rapidly detecting animals in natural scenes. Humans perform this task remarkably well: evoked potential studies indicate that the corresponding neural signals can emerge in the brain within 150 msec of stimulus onset [32], and eye movements toward animal targets can be initiated in roughly the same timeframe [20].

Until recently, little was known about the cues that humans use to achieve this impressive level of performance. However, a recent study by Elder & Velisavljević [10] sheds some light on this question. This study made use of a standard computer vision image dataset called the Berkeley Segmentation Dataset (BSD) [26]. For each image in the dataset, the BSD provides hand segmentations created by human subjects, each of which carves up the image into meaningful regions. Elder & Velisavljević used this dataset to create new images in which luminance, color, texture and shape cues were selectively turned on or off (Fig. 5.1—top). They then measured performance for animal detection using these various modified images over a range of stimulus durations (Fig. 5.1—bottom left), and estimated the weight of each cue using a multiple regression technique (Fig. 5.1—bottom right).

J.H. Elder (✉)
Centre for Vision Research, York University, Toronto, Canada
e-mail: jelder@yorku.ca

S.J. Dickinson, Z. Pizlo (eds.), *Shape Perception in Human and Computer Vision*,
Advances in Computer Vision and Pattern Recognition,
DOI 10.1007/978-1-4471-5195-1_5,

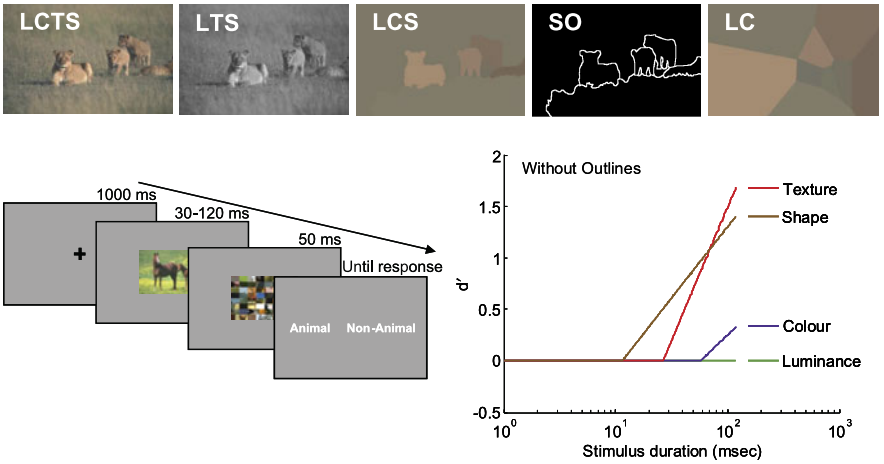

Fig. 5.1 Psychophysical animal detection experiment. *Top*: Example stimuli. The *letters* indicate the cues available: Luminance, Color, Texture, Shape. 'SO' stands for 'Shape Outline'. *Bottom left*: Stimulus sequence. *Bottom right*: Estimated loadings for four cues to animal detection. From [10]

The results show that humans do not use simple luminance or color cues for animal detection, but instead rely on shape and texture cues. Interestingly, shape cues appear to be the first available, influencing performance for stimulus durations as short as 10 msec, within a backward masking paradigm. These results suggest that contour shape cues are not "luxury items" used only when time is not a factor, but rather underlie our fastest judgements about the objects around us. So the question remains: how does the brain rapidly and reliably extract contour shape information from complex natural scenes?

5.2 Computational Models

Computer vision algorithms for contour grouping typically assume as input a map of the local oriented elements to be grouped into chains corresponding to the boundaries of objects in the scene. This is a combinatorial problem—exhaustive methods have exponential complexity and are thus infeasible as algorithms or models for information processing in the brain.

To tame this complexity, most research has focused on modelling and exploiting only the first-order probabilistic relations between successive elements on bounding contours, in other words, modelling contours, either explicitly or implicitly, as first-order Markov chains. In the psychophysics community, this has led to the notion of an "association field" encoding these local relations [13, 28], identified with long-range lateral connections known to link compatible orientation hypercolumms in primate striate cortex [16]. The probabilistic expression of this model has been supported by studies of the ecological statistics of contour grouping, which have also focused principally upon first-order cues [7, 15, 21, 35].

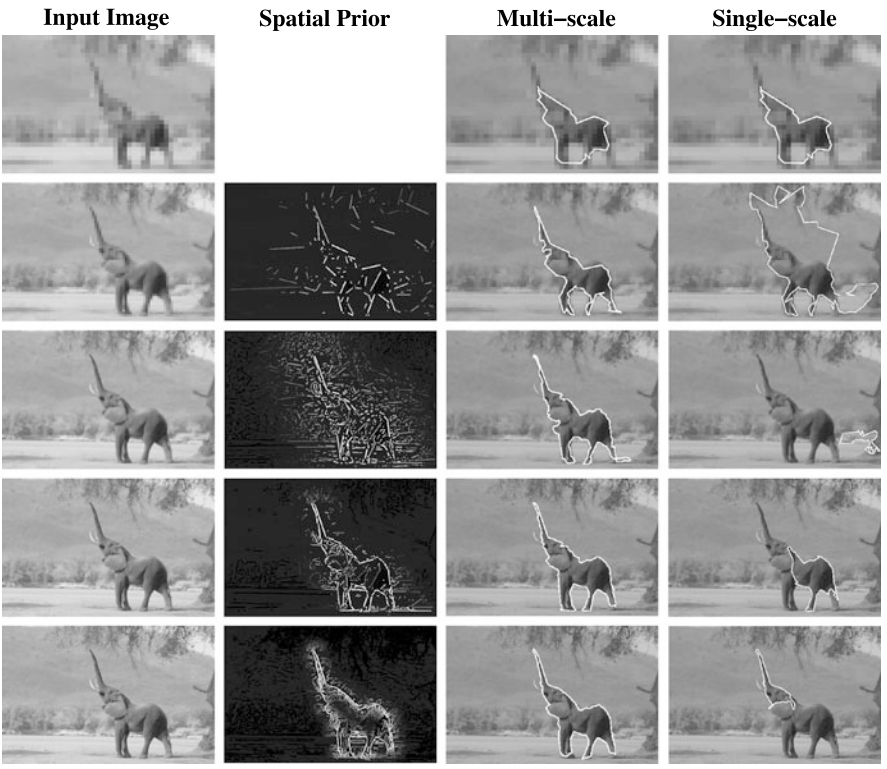

Fig. 5.2 Contour grouping algorithms. *Right column*: Single scale. *Left three columns*: multi-scale, with coarse-to-fine feedback. From [12]

Similarly, many computer vision algorithms for contour grouping have employed a Markov assumption and have focused on first-order cues [3, 4, 8, 11, 18, 24, 28, 33, 41, 43]. However, these first-order Markov algorithms have generally not performed well unless augmented by additional problem-domain knowledge [8] or user interaction [5]. An example from [8] is shown in Fig. 5.2 (right column). The algorithm proceeds by greedy search over the exponential space of possible contours, monotonically increasing the length of the contour hypotheses, and pruning those of lower probability. As can be seen in this example, closed contours corresponding to *parts* of objects can sometimes be computed in this way, but for complex scenes it is rare that the entire object boundary is recovered exactly, unless additional domain-specific constraints are brought to bear.

Recently, however, Estrada & Elder [12] demonstrated that the same algorithm performs much more effectively when placed within a coarse-to-fine scale-space framework (Fig. 5.2—left three columns). In this framework, a Gaussian scale-space over the image is formed, and greedy search is first initiated at the coarsest scale. Since the number of features at this scale is greatly reduced, the search space is much smaller and the algorithm generally finds good, coarse blob hypotheses that code the

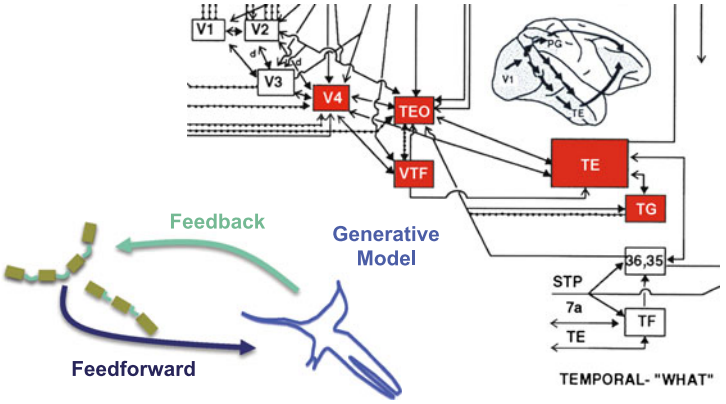

Fig. 5.3 Feedback in the human object pathway. The *diagram on the upper right* is modified from [39]. *Solid arrowheads* indicate feedforward connections, *open arrowheads* indicate feedback connections

rough location and shape of the salient objects in the scene. These hypotheses are then fed back to the next finer level of resolution, where they serve as probabilistic priors, conditioning the likelihoods and effectively shrinking the search space to promising regions of the image.

The success of this simple framework raises the possibility that the brain may also use a feedback mechanism to transmit global shape constraints to the early visual mechanisms involved in contour grouping.

5.3 Feedback in the Primate Object Pathway

Physiologically, it is certainly plausible that feedback might be involved in the perceptual organization of contours in the human brain. Figure 5.3 (right) shows the known connectivity of visual areas in the object pathway of primate brain. While processing is often described by default as a feedforward sequence V1 $\rightarrow$ V2 $\rightarrow$ V4 $\rightarrow$ TE/TEO [37], in fact there are feedback connections from each of the later areas to each of the earlier areas, as well as additional feedforward connections.

While some have argued that animal detection by humans and other primates is too fast to allow time for feedback [20, 32, 37], behavioral and physiological reaction times are always broadly dispersed, with a long positive tail. Thus even if the very fastest reactions (perhaps on the easy conditions) are strictly feedforward, there is time for feedback on the rest. Furthermore, recent evidence suggests that visual signals arrive in higher areas much faster than previously thought [14], allowing sufficient time for feedback even on the fastest trials.

What exactly are the grouping computations effected by the recurrent circuits in primate object pathway? We are far from being able to answer this question, but Fig. 5.3 (left) illustrates one specific conceptual model (see also [2, 22, 38, 42]). For

concreteness, let us suppose that earlier areas (e.g., V1, V2) in the visual pathway compute and encode specific partial grouping hypotheses corresponding to fragments of contours. These fragment hypotheses are communicated to higher-order areas (e.g., V4 or TEO), which use them to generate more complete hypotheses of the global shape. These global hypotheses are then fed back to earlier visual areas to sharpen selectivity for other fragments that might support these global hypotheses.

A central component of this architecture is a generative model of shape that is capable of producing probable global shape hypotheses given partial shape information. This generative shape module will be our focus for the remainder of the chapter.

5.4 Generative Models of Shape

While there are many computational theories and algorithms for shape representation, few are truly generative, and those that are or could be have not been fully developed and tested (e.g., [23]). An important exception is the shapelet theory proposed by Dubinskiy & Zhu [6]. The theory is based upon the representation of a shape by a summation of component *shapelets*. A shapelet is a primitive curve defined by Gabor-like coordinate functions that map arclength to the plane. Shifting and scaling shapelets over the arclength parameter produces a basis set that, when combined additively, can model arbitrarily complex shapes.

The shapelet approach has many advantages. For example, components are localized, albeit only in arclength, and scale is made explicit in a natural way. However, like other contour-based methods, the shapelet theory does not explicitly capture regional properties of shape. Perhaps most crucially, the model does not respect the topology of object boundaries: sampling from the model will in general yield non-simple, i.e., self-intersecting, curves.

5.4.1 Localized Diffeomorphisms: Formlets

A different class of model that could be called region-based involves the application of coordinate transformations of the planar space in which a shape is embedded. This idea can be traced back at least to D'Arcy Thompson, who considered specific classes of global coordinate transformations to model the relationship between the shapes of different animal species [36]. Coordinate transformation methods for representing shape have been explored more recently in the field of computer vision (e.g., [19, 34]), but these methods do not in general preserve the topology of embedded contours.

While general smooth coordinate transformations of the plane will not preserve the topology of an embedded curve, it is possible to design a specific family of diffeomorphic transformations that will [9, 17, 27]. It then follows immediately by induction that a generative model based upon arbitrary sequences of diffeomorphisms will preserve topology.

Specifically, let us consider a family of diffeomorphisms we will call *formlets* [9, 27], in tribute to D'Arcy Thompson's seminal book *On Growth and Form* [36]. A formlet is a simple, isotropic, radial deformation of planar space that is localized within a specified circular region of a selected point in the plane. The family comprises formlets over all locations and spatial scales. While the gain of the deformation is also a free parameter, it is constrained to satisfy a simple criterion that guarantees that the formlet is a diffeomorphism. Since topological changes in an embedded figure can only occur if the deformation mapping is either discontinuous or non-injective, these diffeomorphic deformations are guaranteed to preserve the topology of embedded figures.

This formlet model is closely related to recent work by Grenander and colleagues [17], modeling changes to anatomical parts over time. There the problem is: given two MRI images I_t and I_{t+1} of an anatomical structure taken at two successive times t and $t+1$, first (a) compute the deformation vector field that associates each pixel of I_t with a pixel of I_{t+1}, and then (b) represent this deformation field by a sequence of local and radial diffeomorphisms. They demonstrated their method, which they called Growth by Random Iterated Diffeomorphisms (GRID), on the problem of tracking growth in the rat brain, as revealed in sequential planar sections of MRI data. Subsequent work has focused on the generalization of this method to other coordinate systems [29], on establishing the existence and uniqueness of a continuous 'growth flow' given a specified forcing function [31] and on investigating regularized versions of the GRID formulation [30].

The underlying mathematics here are very similar, although there are some important differences in the exact nature of the localized diffeomorphisms and the manner in which parameters are estimated. But the crucial question of interest here is whether these ideas can be extended to model not just differential deformations between two successive images, but to serve as the framework for a generative model over the entire space of smooth shapes, based upon a universal embryonic shape in the plane such as an ellipse.

5.5 Formlet Coding

5.5.1 Formlet Bases

We represent the image by the complex plane $\mathbb{C}$, and define a formlet $f : \mathbb{C} \to \mathbb{C}$ to be a diffeomorphism localized in scale and space. Such a deformation can be realized by centering f about the point $\zeta \in \mathbb{C}$ and allowing f to deform the plane within a ($\sigma \in \mathbb{R}^+$)-region of ζ. A Gabor-inspired formlet deformation can be defined as

$$\begin{aligned} f(z;\zeta,\sigma,\alpha) &= \zeta + \frac{z-\zeta}{|z-\zeta|}\rho\big(|z-\zeta|;\sigma,\alpha\big), \quad \text{where} \\ \rho(r;\sigma,\alpha) &= r + \alpha \sin\left(\frac{2\pi r}{\sigma}\right)\exp\left(\frac{-r^2}{\sigma^2}\right). \end{aligned} \tag{5.1}$$

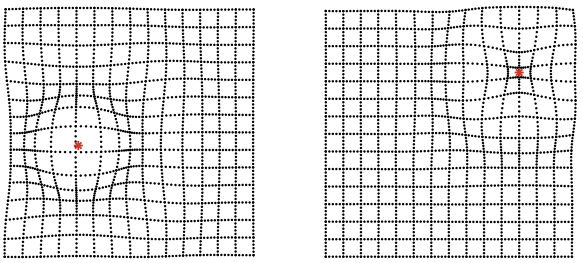

(a) Expansion ($\alpha > 0$) (b) Compression ($\alpha < 0$)

Fig. 5.4 Example formlet deformations. The location ζ of the formlet is indicated by the *red marker*. From [27]

Thus, each formlet $f : \mathbb{C} \to \mathbb{C}$ is a localized isotropic and radial deformation of the plane at location ζ and scale σ. The magnitude of the deformation is controlled by the gain parameter $\alpha \in \mathbb{R}$. Figure 5.4 demonstrates formlet deformations of the plane with positive and negative gain.

5.5.2 Diffeomorphism Constraint

Without any constraints on the parameters, these deformations, though continuous, can fold the plane on itself, changing the topology of an embedded contour. In order to preserve topology, we must constrain the gain parameter to guarantee that each deformation is a diffeomorphism. As the formlets defined in Eq. (5.1) are both isotropic and angle preserving, it is sufficient to require that the radial deformation ρ be a diffeomorphism of $\mathbb{R}^+$, i.e., that $\rho(r; \sigma, \alpha)$ be strictly increasing in r. It can be shown [9, 17, 27] that this requirement leads to a very simple *diffeomorphism constraint*:

$$\alpha \in \sigma\left(-\frac{1}{2\pi}, 0.1956\right). \tag{5.2}$$

By enforcing this constraint, we guarantee that the formlet $f(z, \zeta, \sigma, \alpha)$ is a diffeomorphism of the plane.

Figures 5.5(a) and (b) show the radial deformation function $\rho(r; \sigma, \alpha)$ as a function of r for a range of gain α and scale σ values respectively. Figures 5.5(c) and (d) show the corresponding trace of the formlet deformation of an ellipse in the plane.

5.5.3 Formlet Composition

The power of formlets is that they can be composed to produce complex shapes while preserving topology. Given an embryonic shape $\Gamma^0(t)$ and a sequence of K formlets $\{f_1, \ldots, f_K\}$, the new shape $\Gamma^K(t)$, defined as

$$\Gamma^K(t) = (f_K \circ f_{K-1} \circ \cdots \circ f_1)\big(\Gamma^0(t)\big), \tag{5.3}$$

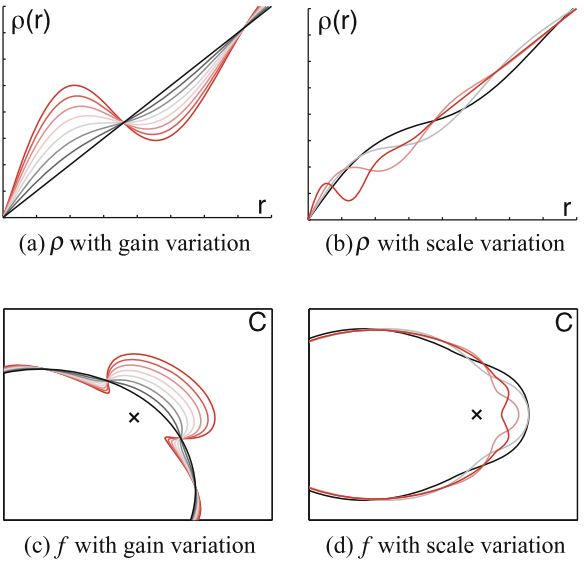

(a) ρ with gain variation (b) ρ with scale variation

(c) f with gain variation (d) f with scale variation

Fig. 5.5 Formlet transformations as a function of scale and gain. *Red* denotes invalid formlet parameters outside the diffeomorphism bounds of Eq. (5.2). From [27]

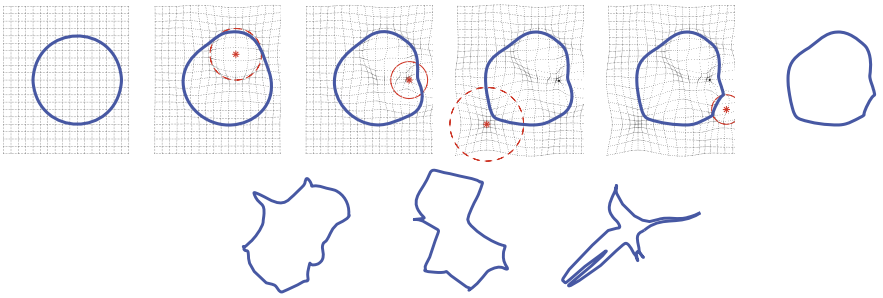

Fig. 5.6 Shapes generated by random formlet composition over the unit circle. *Top row*: shapes resulting from a sequence of five random formlets. The *red dot* and *circle* indicate formlet location ζ and scale σ, respectively. *Bottom row*: example shapes produced from the composition of many random formlets. From [27]

is guaranteed to have the same topology as the original embryonic shape $\Gamma^0(t)$.

Figure 5.6 shows an example of forward composition from a circular embryonic shape, where the formlet parameters ζ, σ, and α have been randomly selected. Note that a rich set of complex shapes is generated without leaving the space of valid shapes (simple, closed contours).

A more difficult but interesting problem is inverse formlet composition: given an *observed* shape $\Gamma^{\text{obs}}(t)$, determine the sequence of K formlets $\{f_1, \ldots, f_K\}$, drawn from a formlet dictionary $\mathcal{D}$ producing the new shape $\Gamma^K(t)$ that best approximates $\Gamma^{\text{obs}}(t)$, according to a specified error measure ξ. Here, we measure error as the L^2 norm of the residual.

5.6 Formlet Pursuit

To explore the inverse problem of constructing formlet representations of planar shapes, Oleskiw et al. [9, 27] employed a set of 391 blue-screened images of animal models from the Hemera Photo-Object database. The boundary of each object was sampled at 128 points at regular arc-length intervals. The full dataset of object shapes used is available at www.elderlab.yorku.ca/formlets.

To estimate the optimal formlet sequence $\{f_1, \ldots, f_K\}$, a version of matching pursuit for sparse approximation was employed [25]. Specifically, given an observed target shape Γ^{obs}, the model was initialized as a an embryonic elliptical shape Γ^0 minimizing the L^2 error $\xi(\Gamma^{\text{obs}}, \Gamma^0)$. At iteration k of the formlet pursuit algorithm, the formlet $f_k(z; \zeta_k, \sigma_k, \alpha_k)$ is selected that, when applied to the current model Γ^{k-1}, maximally reduces the approximation error.

This is a difficult non-convex optimization problem with many local minima. Fortunately, the error function is quadratic in the formlet gain α, so that, given a specified location ζ and scale σ, the optimal gain α^* can be computed analytically [9, 27]. Thus, the problem comes down to a search over location and scale parameters. In practice, this problem can be solved effectively by a *dictionary descent* method, which combines a coarse grid search with local gradient descent at promising locations in the parameter space [9].

5.7 Evaluation

This shape model can be evaluated by addressing the problem of *contour completion*, using the animal shape dataset. In natural scenes, object boundaries are often fragmented by occlusion and loss of contrast: contour completion is the process of filling in the missing parts. Note that this is precisely the task of the generative model in the feedback process illustrated in Fig. 5.3.

Oleskiw and colleagues [9, 27] compared the formlet model with the shapelet model described in Sect. 5.4 [6]. For each shape in the dataset, they simulated the occlusion of a section of the contour, and allowed the two methods to pursue only the remaining visible portion. They then measured the residual error between the model and target for both the visible and occluded portions of the shapes. Performance on the occluded portions, where the model is under-constrained by the data, reveals how well the structure of the model captures properties of natural shapes. Implementations for both the formlet and shapelet models are available at www.elderlab.yorku.ca/formlets.

Figure 5.7 shows example qualitative results for this experiment. While shapelet pursuit introduces topological errors in both visible and occluded regions, formlet pursuit remains topologically valid, as predicted. Figure 5.8 shows quantitative results. While the shapelet and formlet models achieve comparable error on the visible portions of the boundaries, on the occluded portions the error is substantially lower for the formlet representation. This suggests that the structure of the formlet model

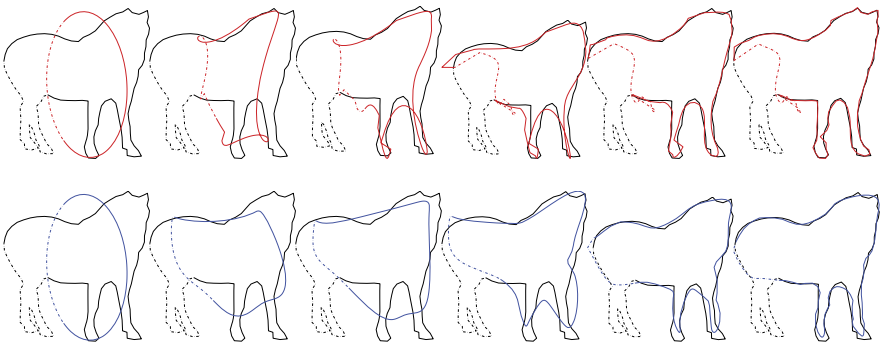

Fig. 5.7 Example of 30 % occlusion pursuit with shapelets (*red*) and formlets (*blue*) for $k = 0, 2, 4, 8, 16, 32$. *Solid lines* indicate visible contour, *dashed lines* indicate occluded contour. From [9]

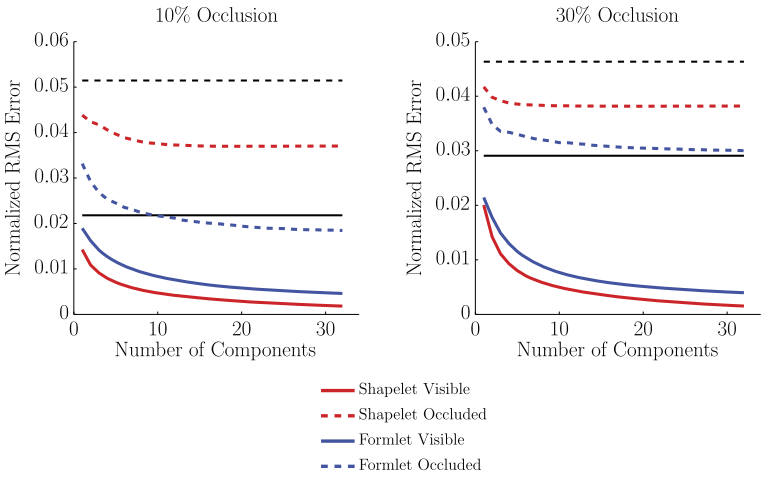

Fig. 5.8 Results of occlusion pursuit evaluation. *Black* denotes error for $\Gamma^0(t)$, the affine-fit ellipse. From [9]

better captures regularities in the shapes of natural objects. The two principal reasons for this are thought to be [9] (a) respecting the topology of the shape prunes off many inferior completion solutions and (b) by working in the image space, rather than arc length, the formlet model is better able to capture important regional properties of shape.

5.8 Discussion

Strictly feedforward algorithms for contour grouping based upon first-order Markov models of contours tend to work poorly on complex natural scenes, yet humans

are able to make effective use of contour shape information for object detection [10]. Proven performance advantages of coarse-to-fine methods for contour grouping [12], together with the massive feedback connections that are known to pervade primate object pathway [39, 40] suggest that the human brain may employ a recurrent computation to group contours and efficiently extract shape information from natural scenes.

A key requirement for this recurrent network is a generative model of shape capable of producing global shape "hallucinations" based on contour fragments computed in early visual cortex. These global shape hypotheses can then be fed back to early visual areas to condition search for additional fragments that might support the hypotheses.

The main problem in establishing such a generative model has been topology: prior models do not guarantee that sampled shapes are simple closed contours (e.g., [6]). Recently, however, a novel framework for shape representation has been introduced that guarantees that sampled shapes will have the correct topology. The theory [9, 27], based upon localized diffeomorphic deformations of the image called *formlets*, has its roots in early investigations of biological shape transformation [36], and is closely related to recent work modelling growth in biomedical imaging data [17]. The formlet representation is seen to yield more accurate shape completion than an alternative contour-based generative model of shape, which should make it more effective at generating global shape hypotheses to guide feedforward contour grouping processes.

These findings suggest a number of future experiments and computational investigations: (1) Is there any psychophysical evidence that humans exploit higher-order shape features to segment contours in complex, cluttered scenes? If we do, is there any evidence that this involves a feedback circuit? (2) Many shapes have highly elongated parts that are not efficiently modelled by isotropic formlets. Is there a way to generalize the theory to incorporate oriented formlets? (3) Applying the theory effectively for problems of grouping, detection and recognition will require a probabilistic model over formlet sequences. What is an appropriate structure for this model, and how can its parameters be learned?

Answers to these questions will bring us considerably closer to an understanding of the perceptual organization of shape.

Acknowledgements The author would like to thank Francisco Estrada, Tim Oleskiw, Gabriel Peyré, Ljiljana Velisavljević and Alexander Yakubovich for their contributions to the work reviewed in this chapter. This work was supported by NSERC, OCE and GEOIDE.

References

1. Arbelaez P, Maire M, Fowlkes C, Malik J (2011) Contour detection and hierarchical image segmentation. IEEE Trans Pattern Anal Mach Intell 33(5):898–916. doi:10.1109/TPAMI.2010.161
2. Cavanagh P (1991) What's up in top-down processing. In: Gorea A (ed) Representations of vision. Cambridge University Press, Cambridge, pp 295–304

3. Cohen LD, Deschamps T (2001) Multiple contour finding and perceptual grouping as a set of energy minimizing paths. In: Energy minimization methods workshop, CVPR
4. Crevier D (1999) A probabilistic method for extracting chains of collinear segments. Comput Vis Image Underst 76(1):36–53
5. Deschamps T, Cohen LD (2000) Minimal paths in 3d images and application to virtual endoscopy. In: Proceedings of the 6th European conference on computer vision
6. Dubinskiy A, Zhu SC (2003) A multiscale generative model for animate shapes and parts. In: Proc. 9th IEEE ICCV, vol 1, pp 249–256
7. Elder JH, Goldberg RM (2002) Ecological statistics of Gestalt laws for the perceptual organization of contours. J Vis 2(4):324–353
8. Elder JH, Krupnik A, Johnston LA (2003) Contour grouping with prior models. IEEE Trans Pattern Anal Mach Intell 25(25):661–674
9. Elder JH, Oleskiw TD, Yakubovich A, Peyré G (2013) On growth and formlets: sparse multiscale coding of planar shape. Image Vis Comput 31:1–13
10. Elder JH, Velisavljević L (2009) Cue dynamics underlying rapid detection of animals in natural scenes. J Vis 9(7):1–20
11. Elder JH, Zucker SW (1996) Computing contour closure. In: Proceedings of the 4th European conference on computer vision. Springer, New York, pp 399–412
12. Estrada FJ, Elder JH (2006) Multi-scale contour extraction based on natural image statistics. In: IEEE conference on computer vision and pattern recognition workshop
13. Field DJ, Hayes A, Hess RF (1993) Contour integration by the human visual system: evidence for a local "association field". Vis Res 33(2):173–193
14. Foxe JJ, Simpson GV (2002) Flow of activation from V1 to frontal cortex in humans. Exp Brain Res 142:139–150
15. Geisler WS, Perry JS, Super BJ, Gallogly DP (2001) Edge co-occurence in natural images predicts contour grouping performance. Vis Res 41(6):711–724
16. Gilbert CD, Wiesel TN (1989) Columnar specificity of intrinsic horizontal and corticocortical connections in cat visual cortex. J Neurosci 9(7):2432–2443
17. Grenander U, Srivastava A, Saini S (2007) A pattern-theoretic characterization of biological growth. IEEE Trans Med Imaging 26(5):648–659
18. Cox IJ, Rehg JM, Hingorani S (1993) A Bayesian multiple hypothesis approach to contour segmentation. Int J Comput Vis 11(1):5–24
19. Jain AK, Zhong Y, Lakshmanan S (1996) Object matching using deformable templates. IEEE Trans Pattern Anal Mach Intell 18(3):267–278
20. Kirchner H, Thorpe SJ (2006) Ultra-rapid object detection with saccadic eye movements: visual processing speed revisited. Vis Res 46:1762–1766
21. Kruger N (1998) Collinearity and parallelism are statistically significant second order relations of complex cell responses. Neural Process Lett 8:117–129
22. Lee TS, Mumford D (2003) Hierarchical Bayesian inference in the visual cortex. J Opt Soc Am A 20(7):1434–1448
23. Leyton M (1988) A process-grammar for shape. Artif Intell 34(2):213–247
24. Mahamud S, Thornber KK, Williams LR (1999) Segmentation of salient closed contours from real images. In: IEEE international conference on computer vision. IEEE Computer Society, Los Alamitos, pp 891–897
25. Mallat SG, Zhang Z (1993) Matching pursuits with time frequency dictionaries. IEEE Trans Signal Process 41(12):3397–3415
26. Martin D, Fowlkes C, Malik J (2004) Learning to detect natural image boundaries using local brightness, color and texture cues. IEEE Trans Pattern Anal Mach Intell 26(5):530–549
27. Oleskiw TD, Elder JH, Peyré G (2010) On growth and formlets. In: Proceedings of the IEEE conference on computer vision and pattern recognition (CVPR)
28. Parent P, Zucker SW (1989) Trace inference, curvature consistency, and curve detection. IEEE Trans Pattern Anal Mach Intell 11:823–839
29. Portman N, Grenander U, Vrscay ER (2009) Direct estimation of biological growth properties from image data using the "GRID" model. In: Kamel M, Campilho A (eds) Image analysis

and recognition, proceedings. Lecture notes in computer science, vol 5627, pp 832–843. 6th international conference on image analysis and recognition, Halifax, Canada, Jul 06–08, 2009
30. Portman N, Grenander U, Vrscay ER (2011) GRID macroscopic growth law and it application to image inference. Q Appl Math 69(2):227–260
31. Portman N, Vrscay ER (2011) Existence and uniqueness of solutions to the grid macroscopic growth equation. Appl Math Comput 217(21):8318–8327. doi:10.1016/j.amc.2011.03.021. http://www.sciencedirect.com/science/article/pii/S0096300311003754
32. Thorpe S, Fize D, Marlot C (1996) Speed of processing in the human visual system. Nature 381:520–522
33. Sha'ashua A, Ullman S (1988) Structural saliency: the detection of globally salient structures using a locally connected network. In: Proceedings of the 2nd international conference on computer vision, pp 321–327
34. Sharon E, Mumford D (2004) 2d-shape analysis using conformal mapping. In: Computer vision and pattern recognition, IEEE Comp. Soc. Conf., vol 2, pp 350–357
35. Sigman M, Cecchi GA, Gilbert CD, Magnasco MO (2001) On a common circle: natural scenes and Gestalt rules. Proc Natl Acad Sci 98(4):1935–1940
36. Thompson DW (1961) On growth and form. Cambridge University Press, Cambridge
37. Thorpe S (2002) Ultra-rapid scene categorization with a wave of spikes. In: Bülthoff HH et al (eds) Proceedings of the biologicaly-motivated computer vision conference. LNCS, vol 2525, pp 1–15
38. Tu Z, Chen X, Yuille AL, Zhu SC (2005) Image parsing: unifying segmentation, detection, and recognition. Int J Comput Vis 63(2):113–140
39. Ungerleider L (1995) Functional brain imaging studies of cortical mechanisms for memory. Science 270(5237):769–775
40. Van Essen DC, Olshausen B, Anderson CH, Gallant JL (1991) Pattern recognition, attention, and information processing bottlenecks in the primate visual search. SPIE 1473:17–28
41. Williams LR, Jacobs DW (1997) Stochastic completion fields: a neural model of illusory contour shape and salience. Neural Comput 9(4):837–858
42. Yuille A, Kersten D (2006) Vision as Bayesian inference: analysis by synthesis? Trends Cogn Sci 10(7):301–308
43. Zucker SW, Hummel R, Rosenfeld A (1977) An application of relaxation labeling to line and curve enhancement. IEEE Trans Comput 26:394–403