



HAL
open science

On Growth and Formlets: Sparse Multi-Scale Coding of Planar Shape

Timothy D. Oleskiw, James Elder, Gabriel Peyré

► **To cite this version:**

Timothy D. Oleskiw, James Elder, Gabriel Peyré. On Growth and Formlets: Sparse Multi-Scale Coding of Planar Shape. CVPR'10, Jun 2010, San-Francisco, United States. pp.459-466 . hal-00470560

HAL Id: hal-00470560

<https://hal.science/hal-00470560>

Submitted on 6 Apr 2010

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

On Growth and Formlets: Sparse Multi-Scale Coding of Planar Shape

Timothy D. Oleskiw
York University
oleskiw@yorku.ca

James H. Elder
York University
jelder@yorku.ca

Gabriel Peyré
Université Paris-Dauphine
gabriel.peyre@ceremade.dauphine.fr

Abstract

This paper presents a sparse representation of 2D planar shape through the composition of warping functions, termed formlets, localized in scale and space. Each formlet subjects the 2D space in which the shape is embedded to a localized isotropic radial deformation. By constraining these localized warping transformations to be diffeomorphisms, the topology of shape is preserved, and the set of simple closed curves is closed under any sequence of these warpings. A generative model based on a composition of formlets applied to an embryonic shape, e.g., an ellipse, has the advantage of synthesizing only those shapes that could correspond to the boundaries of physical objects. To compute the set of formlets that represent a given boundary, we demonstrate a greedy coarse-to-fine formlet pursuit algorithm that serves as a non-commutative generalization of matching pursuit for sparse approximations. We evaluate our method by pursuing partially occluded shapes, comparing performance against a contour-based sparse shape coding framework.

1. Introduction

Shape information is important for a broad range of computer vision problems. For some detection and recognition tasks, discriminative models that use non-invertible shape codes [2] can be effective. However, many other tasks call for a more complete generative model of shape. Examples include: (1) shape segmentation, recognition, and tracking in cluttered scenes, where shapes must be distinguished not just from each other, but from ‘phantom’ shapes formed by conjunctions of features from multiple objects [6]; (2) modeling of shape articulation, growth, and deformation; and (3) modeling of shape similarity.

Our paper concerns the generative modeling of natural 2D shapes in the plane, represented by their 1D boundary. We restrict our attention to simply-connected shapes whose boundaries are smooth, simple, and closed curves. We seek a generative shape model that satisfies a set of properties that seem to us essential:

1. *Completeness.* The model can produce all shapes.
2. *Closure.* The set of valid shapes is closed under the generative model. In other words, the model generates only valid shapes.
3. *Composition.* Complex shapes are generated by combining simpler components.
4. *Sparsity.* Good approximations of shape can be generated with relatively few components.
5. *Progression.* Approximations can be improved by incorporating more components.
6. *Locality.* Components are localized in space.
7. *Scaling.* Components are tuned to specific scales and are self-similar over scale.
8. *Region & Contour.* Components can capture both region and contour properties in a natural way.

The need for completeness is self-evident if the system is to be general. Closure is critical if we hope to capture the statistics of natural shape in a set of hidden generative variables. Without closure, heuristics must be used to avoid the generation of invalid shapes, e.g., bounding contours with self-intersections. Aside from the resulting inefficiency, this creates a discrepancy between the statistical structure encoded by the model, and samples the model produces. In other words, the model cannot fully capture the statistics of natural boundaries.

Composition (here we use the word in a general sense) is important if we are to handle the richness and complexity of natural shapes while maintaining conceptual simplicity. Given the high dimensionality of natural shapes, sparsity is necessary in order to store shape models [16]. Sparsity also implies that essential shape features have been made explicit [1]. Progression allows the complexity of the model to be matched to the difficulty of the task, facilitating real-time operation and coarse-to-fine optimization.

Locality is a natural goal, since a first-order property of natural images is local coherence. Nearby points on the surface of an object tend to have similar reflectance, attitude, and illumination. Locality also allows for greater robustness to occlusion, since components are more likely to be either entirely visible or removed altogether rather than distorted.

Scaling allows invariance over object size, and allows shape features of different sizes to be captured separately.

Finally, it has long been recognized that planar shape description requires attention to both region and contour properties [16]. Some shape properties, *e.g.*, curvature, are naturally described by the bounding contour. Others, *e.g.*, necks, are best described as region properties, since they involve points that are close together in space but distant along the contour.

A good generative model will allow both to be encoded in a natural way. We begin by reviewing prior models, with an eye to each of these essential properties.

2. Prior Work

Early models that used chain coding or splines to encode shapes were not generative and failed to succinctly capture global properties of shape. Fourier descriptor, moment, and PCA bases have the potential to be generative, but since all components are global, they are not robust to occlusion or local deformation [7, 16, 18]. For these reasons, most modern approaches attempt to capture structure at intermediate scales, or over a range of scales. Most of these models can be crudely partitioned into two classes: contour-based and symmetry-based.

2.1. Contour-Based Models

Attneave [1] pointed to the concentration of information in the curvature of the bounding contour, and suggested the potential for sparse descriptions based on points of extremal curvature magnitude. Hoffman & Richards [10] linked curvature to the part structure of shapes, proposing that parts are perceptually segmented at negative minima of curvature. Mokhtarian & colleagues emphasized the encoding of curvature inflections across scale space for the purpose of shape recognition [15].

While none of these early models are generative, Dubinskiy & Zhu [8] have more recently proposed a contour-based shape representation that is both generative and sparse. 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 in arclength, and scale is made explicit in a natural way. However, like all 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.

2.2. Symmetry-Based Models

Blum and colleagues [3, 4] introduced the symmetry axis representation of shape in which a planar shape is represented by a 1D skeleton function and associated 1D radius function. The symmetry axis representation led to related representations [5] which found application in medical imaging and other domains.

Subsequent work incorporated notions of scale and time with symmetry axis descriptions. Leyton [13] related symmetry axis descriptions to causal deformation processes acting upon prototype shapes. In this view, symmetry axes, terminating at curvature extrema on the boundary, are understood as records of these deformation processes. Subsequent work on curve evolution methods and shock-graph representations [12, 17] has provided a more complete theory of region-based shape representations that have been broadly applied.

Despite the many appealing features of symmetry axis and shock-graph representations, these methods, in general, are not sparse. In fact, the description of each shape typically requires more storage, and little emphasis has been placed on making symmetry axis representations generative [16]. Very recent work of Trinh and Kimia exploring generative and sparse models based upon shock graphs comes some way in overcoming these limitations [21]. However, the constraints required to enforce the closure property, *i.e.*, topological constraints, are fairly complex, and the full potential of the theory has yet to be explored.

2.3. Hybrid Approaches

Recognizing the merits and limitations of both contour-based and symmetry-based approaches, Zhu [22] developed an MRF model for natural 2D shape, employing a neighbourhood structure that can directly encode both contour-based and region-based Gestalt principles. The theory is promising in many respects. It is generative, providing an explicit probabilistic model, and it captures both region and contour properties. It is not sparse, however, and because the underlying graph is lifted from the image plane, there is nothing in the model that encodes the topological constraint that the boundary be simple, *i.e.*, non-intersecting. Instead, when sampling from the model, a ‘firewall’ is employed to prevent intersections. Again, this is inefficient, and it also creates a disconnect between the generative variables encoding the model and the sampling distribution.

2.4. Coordinate Transformations

A different class of model that could also 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 of the plane to model the relationship between the shapes of different animal species [20]. In the field of computer vision, Jain *et al.* [11] were among the first to extend this idea to more general deformations with a complete Fourier deformation basis that they used to match observed shapes to stored prototypes. However, this Fourier basis fails to satisfy the locality property, and as a potential generative model it does not satisfy the closure property: random combinations of Fourier deformation components will not in general preserve the topology of the prototype curve.

More recently, Sharon & Mumford [19] have explored conformal mappings as global coordinate transformations between planar shapes. However, although the Riemann mapping theorem guarantees that any simple closed curve can be conformally mapped to the unit circle, conformal mappings do not in general preserve the topology of embedded contours. Hence, despite the computational constraints imposed by the Cauchy-Riemann equations, we again have the problem that the set of valid bounding contours is not closed under these transformations, making generative modeling difficult.

2.5. Localized Diffeomorphisms: Formlets

In considering prior generative shape models, the goal that seems most elusive is that of closure: ensuring that the model generates only valid shapes. Our approach originates with the observation that, while general smooth coordinate transformations of the plane will not preserve the topology of an embedded curve, it is straightforward to design a specific family of diffeomorphic transformations that will. It then follows immediately by induction that a generative model based upon arbitrary sequences of diffeomorphisms will satisfy the closure property.

In this paper we specifically consider a family of diffeomorphisms we call *formlets*. 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. Thus the model satisfies the closure property.

By construction, formlets satisfy the desired locality and scaling properties. It is easy to show that the model also satisfies the composition, completeness, and progression properties in that an arbitrary shape can be approximated to increasing precision by composing an appropriate sequence of localized formlets. Since each formlet may be centered

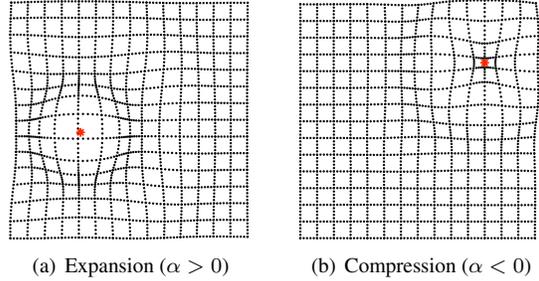


Figure 1. Formlet transformations of a planar grid.

either near the contour, near a symmetry axis, or at any other location in the plane, the model has the potential to capture both region and contour properties directly.

Our formlet model is in part inspired by recent work by Grenander *et al.* [9], modeling changes to anatomical parts over time. Their representation, called Growth by Random Iterated Diffeomorphisms (GRID), models growth as a sequence of local and radial deformations. They demonstrate their model by tracking growth in the rat brain, as revealed in sequential planar sections of MRI data.

In the present paper we explore the possibility that these ideas could be extended to model not just differential growth between sequential shapes, but to serve as the basis for a generative model over the entire space of smooth shapes, based upon a universal embryonic shape in the plane, *i.e.*, an ellipse. Our specific contributions are:

1. We illustrate the completeness and closure properties of the formlet model through random generation of sample shapes.
2. To solve the inverse problem of modeling given shapes, we develop and apply a generalization of matching pursuit, which selects the sequence of formlets that minimizes approximation error. We demonstrate that this *formlet pursuit* algorithm allows for an efficient and progressive approximation of shape, while preserving topological properties.
3. We assess the robustness of the formlet model to occlusion by evaluating it on the problem of contour completion. We find that the model compares favorably with the contour-based shapelet model [8] on this problem.

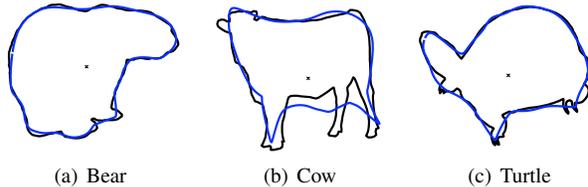


Figure 2. Sparse formlet reconstruction of canonical animal shape.

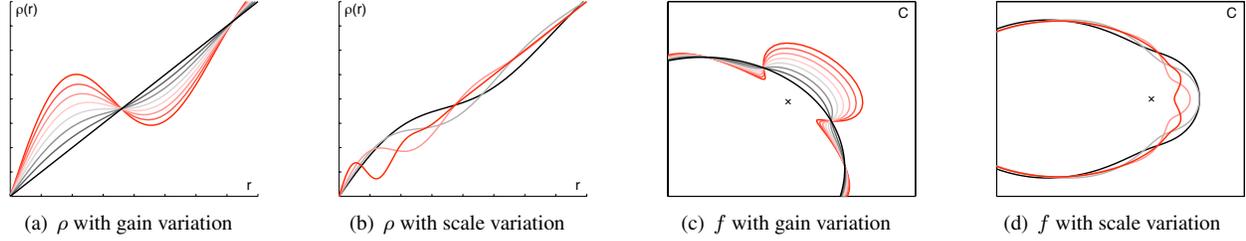


Figure 3. Gabor mapping transformations as scale and gain vary (a,b) with proportional formlets applied to an ellipse (c,d) at a fixed location in space. Red denotes formlet parameters outside the monotonicity bounds of Equation 2.

3. Formlet Coding

3.1. Formlet Bases

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

$$f(z; \zeta, \sigma, \alpha) = \zeta + \frac{z - \zeta}{|z - \zeta|} \rho(|z - \zeta|; \sigma, \alpha), \quad \text{where} \quad (1)$$

$$\rho(r; \sigma, \alpha) = r + \alpha \sin\left(\frac{2\pi r}{\sigma}\right) \exp\left(\frac{-r^2}{\sigma^2}\right).$$

Such formlets are isotropic and radial deformations of the plane with scale σ located at ζ . The severity of deformation is controlled by the gain parameter $\alpha \in \mathbb{R}$. Figure 1 demonstrates formlet deformations of the plane with positive and negative gain.

Without any constraints on the parameters, these deformations, though continuous, can fold the plane on itself, changing the topology of an embedded contour. However, we show in Appendix B that if the scale σ and gain α satisfy

$$-(2\pi)^{-1} < \frac{\alpha}{\sigma} \lesssim 0.1956, \quad (2)$$

the radial deformation function $\rho(r; \sigma, \alpha)$ will be strictly increasing in r , and the formlet $f(z; \zeta, \sigma, \alpha)$ will be diffeomorphic. Hence, such a formlet acting on a curve embedded in the plane will be a homeomorphism between curve topologies. In particular, let Γ be the continuous mapping

$$\Gamma : [0, 1] \in \mathbb{R} \rightarrow \mathbb{C}. \quad (3)$$

Recall that Γ is simple if the mapping is injective, and closed by permitting the equality $\Gamma(0) = \Gamma(1)$. Since a formlet f satisfying Equation 2 is bicontinuous, it is easy to see that if Γ is simple and closed, the deformed curve

$$\Gamma^f(t) = f(\Gamma(t)) \quad (4)$$

will also be simple and closed.

Figures 3(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 3(c) and (d) show the corresponding formlet deformation of an ellipse.

A dictionary \mathcal{D} of formlet transformations can be created by sampling over the range of valid location, scale, and gain parameters, *i.e.*, $(\zeta, \sigma, \alpha) \in \mathbb{C} \times \mathbb{R}^+ \times \mathbb{R}$.

3.2. Formlet Composition: Forward and Inverse

We define the *forward formlet composition* problem as follows. Given an embryonic shape $\Gamma^0(t)$ and a sequence of K formlets $\{f_1 \dots f_K\}$ drawn from a formlet dictionary \mathcal{D} , determine the resulting deformed shape $\Gamma^K(t)$. The problem is well-posed because the set of simple closed curves is closed under formlet deformation: multiple formlets can be composed to generate complex shape transformations. Thus,

$$\Gamma^K(t) = (f_K \circ f_{K-1} \circ \dots \circ f_1)(\Gamma^0(t)). \quad (5)$$

Figure 4 shows an example of forward composition where the formlet parameters ζ, σ , and α have been randomly selected. Note that a complex shape is generated without introducing topological error.

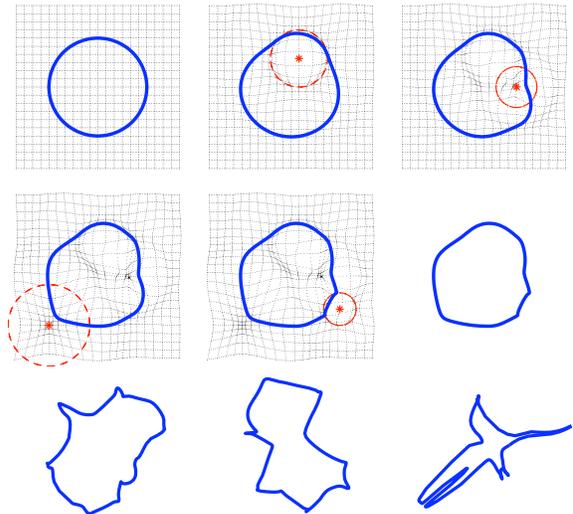


Figure 4. Shapes generated by random formlet composition over the unit circle. Red denotes formlet location and relative scale.

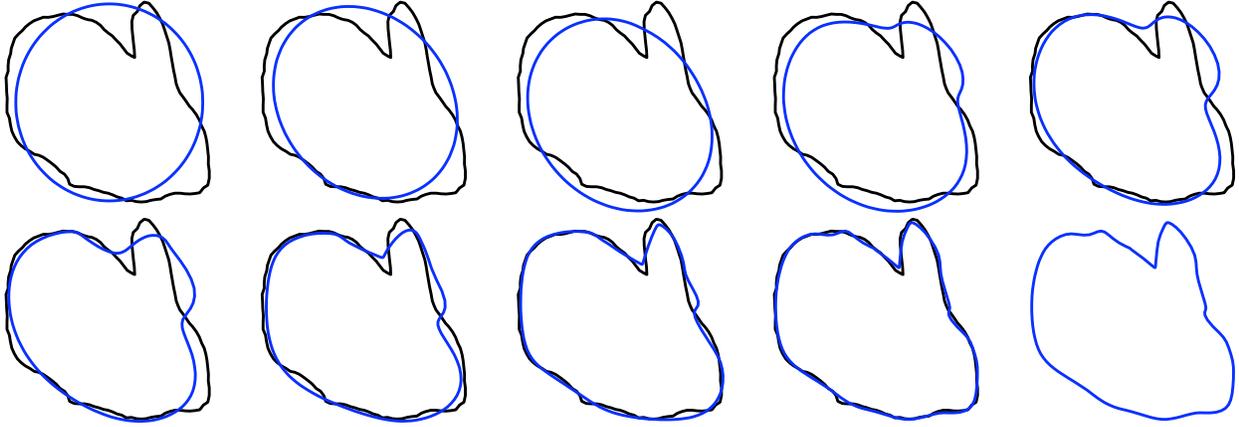


Figure 5. Formlet pursuit of a normalized rabbit shape. Pursuit begins with the unit circle, followed by affine fit, then $K=1,2,3,4,8,16,32$.

A more difficult but interesting problem is *inverse formlet composition*: given an *observed* shape $\Gamma^{obs}(t)$, determine a sequence of K formlets $\{f_1 \dots f_K\}$, drawn from a formlet dictionary \mathcal{D} , that best approximate $\Gamma^{obs}(t)$ by minimizing some reconstruction error ξ . Here we measure error as the L_2 norm of the residual:

$$\begin{aligned} \xi(\Gamma^{obs}, \Gamma^K) &= \|\Gamma^{obs}(t) - \Gamma^K(t)\|_2 \\ &= \int_0^1 (\Gamma^{obs}(t) - \Gamma^K(t)) \overline{(\Gamma^{obs}(t) - \Gamma^K(t))} dt. \end{aligned}$$

The inverse problem is then to find the sequence of formlets $\{f_1 \dots f_K\}$ that minimizes $\xi(\Gamma^{obs}, \Gamma^K)$.

4. Formlet Pursuit

To estimate the optimal formlet sequence $\{f_1 \dots f_K\}$, we propose a version of matching pursuit for sparse approximation [14], replacing the linear summation of elements by a non-commutative composition of formlet components. Figure 2 shows an example of progressive approximation by formlet pursuit. Algorithm 1 shows the flow of the formlet pursuit algorithm.

Initialization Given an observed shape Γ^{obs} , we first regularly sample the contour at constant speed and normalize to satisfy

$$\int_0^1 \mathbf{Re}(\Gamma^{obs}(t))^2 dt = \int_0^1 \mathbf{Im}(\Gamma^{obs}(t))^2 dt = 1. \quad (6)$$

We then initialize formlet pursuit by fitting the observed shape with an ellipse, *i.e.*, by applying an affine transformation $A\Gamma_0 + z_0$ to the unit circle Γ_0 . The mapping $A : \mathbb{C} \rightarrow \mathbb{C}$ is computed deterministically as the solution to a linear system minimizing the error $\xi(\Gamma^{obs}, A\Gamma_0 + z_0)$ with z_0 being the center mass of Γ^{obs} .

Formlet Selection At iteration k of the formlet pursuit algorithm, we select the formlet $f_k(z; \zeta_k, \sigma_k, \alpha_k)$ that, when

applied to the current approximation Γ^{k-1} , maximally reduces the approximation error:

$$f_k = \operatorname{argmin}_{f \in \mathcal{D}} \xi(\Gamma^{obs}, f(\Gamma^{k-1})). \quad (7)$$

This is a difficult non-convex optimization problem, and experimentation with gradient descent methods has shown that the formlet parameter space has many local minima. Fortunately, solving for f_k can be simplified when a formlet transformation is linear with respect to α . Specifically, one can solve for the optimal error-reducing gain α_k when the other parameters are known, similar to the way Dubinskiy *et al.* solve for an optimal shapelet affine transformation [8]. Expressing a formlet defined in Equation 1 as

$$\begin{aligned} f(z; \zeta, \sigma, \alpha) &= z + \alpha \cdot g(z - \zeta, \sigma) \quad \text{where} \\ g(z_\zeta; \sigma) &= \frac{z_\zeta}{|z_\zeta|} \sin\left(\frac{2\pi|z_\zeta|}{\sigma}\right) \exp\left(-\frac{|z_\zeta|^2}{\sigma^2}\right), \end{aligned} \quad (8)$$

the optimal gain α_k is shown in Appendix C to be

$$\alpha_k = \frac{\int_0^1 \langle \Gamma^{obs}(t) - \Gamma^{k-1}(t), g(\Gamma^{k-1}(t) - \zeta_k; \sigma_k) \rangle dt}{\|g(\Gamma^{k-1}(t) - \zeta_k; \sigma_k)\|_2} \quad (9)$$

where $\langle \cdot, \cdot \rangle$ is the standard inner product on \mathbb{C} .

Algorithm 1: Formlet Pursuit of Γ^{obs} .

Initialization: define $\Gamma^0 = A\Gamma_0 + z_0$ to be a best matching ellipse approximating Γ^{obs}

for $k = 1, \dots, K$ **do**

Optimal Formlet: compute maximal error reducing transformation

$$f_k = \operatorname{argmin}_{f \in \mathcal{D}} \xi(\Gamma^{obs}, f(\Gamma^{k-1}))$$

Update Approximation: apply optimal formlet

$$\Gamma^k = f_k(\Gamma^{k-1})$$

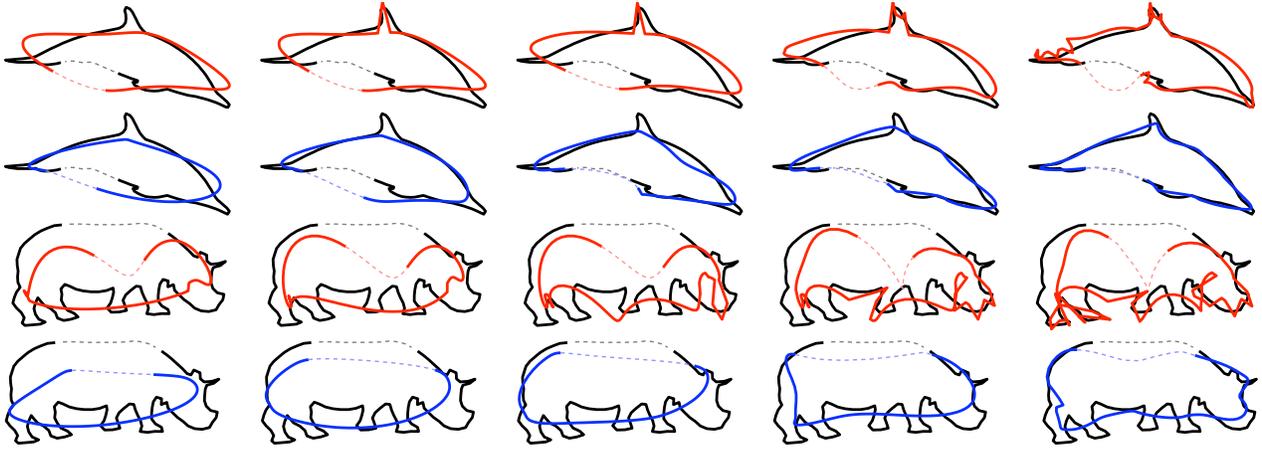


Figure 6. Ocluded pursuit of dolphin and rhino shapes with shapelets (red) and formlets (blue) for $K = 1, 2, 4, 8, 16$. Dashed lines denote 10% occlusion interval. Note the large shape discrepancy of occluded dolphin body and topological errors of shapelet rhino approximation.

Our approach in this paper is therefore to exhaustively search over a regular sampling of the location ζ and scale σ parameters and then solve for the associated optimal gain α by Equation 9. We note our method could likely be made more efficient by a coarse sampling of the parameter space followed by selective gradient descent.

5. Evaluation

We evaluated the proposed formlet representation on a database of natural object shapes, comparing against the multi-scale contour-based *shapelet* representation proposed by Dubinskiy and Zhu [8].

Specifically, we evaluated these two methods on a *occluded pursuit* problem in which silhouette shapes of objects are partially occluded. The algorithms were provided with only the visible portion of target shape, and were thus required to hallucinate the occluded portion. We measured the residual error between the approximation and target shape 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 the degree of consistency between the structure of the model and typical natural shapes.

The database consisted of 391 blue-screened images of animal models from the Hemera object database. The boundary of each object was sampled at regular arc-length intervals: 128 samples per object were used for our experiments. Target shapes were occluded over a randomly-selected contiguous section of their boundaries comprising 10% or 30% of total arclength.

Both the formlet and shapelet representations rely upon sampling certain parameters; we attempted to make the comparison fair by sampling as finely as was feasible given time constraints. The shapelet representation assumes an arc-length representation of the curves on $t \in [0, 1]$, and each shapelet component has an arc-length position μ and scale σ . We sampled the position parameter μ at 128

regularly-spaced values over $[0, 1]$, and the scale parameter σ at 128 regularly-spaced values over $(0, \frac{1}{3}]$ (scales over $1/3$ cause significant topological errors in the resulting shapelets). Additional optimal affine parameters were computed directly [8].

For the formlet approach, the position parameter ζ was sampled on a 128×128 grid roughly 4 times the extent of the average shape, and the scale parameter σ at 128 regularly-spaced values over $(0, 1]$. The optimal gain parameter α was computed directly as described in Section 4.

Both formlet and shapelet pursuit were governed by a minimization of the L_2 error over the visible portion of the curves only. Both shapelet and formlet models were initialized with the same ellipse approximations. We measure performance using a generalization of the Hausdorff distance. Specifically, we define the error between the target shape and the approximation as the average minimum distance of a point on one of the shapes to the other. See Appendix A for supplementary information.

5.1. Results

Figure 6 shows some example qualitative results from this experiment. Note that as pursuit progresses, topological errors are induced in the shapelet representation at both the visible and occluded regions, whereas the formlet model remains topologically correct.

Figure 7 shows the quantitative results of this experiment. A first observation is that, for a fixed budget of components, the shapelet and formlet models achieve comparable error on the visible portions of the boundaries. However, we see a different pattern for the occluded portions: for a given component budget, approximation error is substantially less for the formlet representation. The result is that for a given visible error, the occluded error is substantially lower for formlet approximations. This suggests that by preserving the topology of shape throughout pursuit, formlet coding also better exhibits quantitative regularities present in the shapes of natural objects.

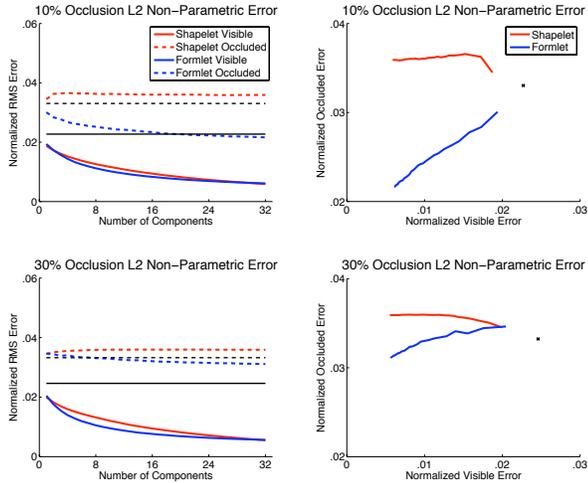


Figure 7. Results of occlusion pursuit with occlusion ratios of 10% and 30%. Black denotes error of the affine-fit ellipse.

6. Conclusion

We have developed a novel generative model of planar shape that satisfies a number of essential properties. In this model, complex shapes are seen as the evolution of a simple embryonic shape by successive application of simple diffeomorphic transformations of the plane called *formlets*. The system is both complete and closed, since arbitrary shapes can be modeled, and generated shapes are guaranteed to be topologically valid. This means that the model has the potential to support accurate probabilistic modeling. We have demonstrated a novel formlet pursuit algorithm that selects formlets to efficiently approximate given target shapes. Evaluation of the formlet pursuit model on the problem of shape completion revealed that the model is better able to approximate parts of shapes missing due to occlusion than a competing contour-based method.

Future Work We hope to extend the present work in a number of ways. First, we would like to increase the expressiveness of formlet components by replacing the scalar gain parameter with an affine transformation. This will allow the model to more efficiently estimate limbs and other elongated features. A second objective is to increase the efficiency of formlet pursuit by combining a coarsely quantized formlet dictionary with selective gradient descent. Finally, we are interested in probabilistic modeling of shapelet components and their relationships, which may lead to applications in recognition, synthesis, and denoising.

References

[1] F. Attneave. Some informational aspects of visual perception. *Psychol. Rev.*, 61(3):183–193, May 1954.
 [2] S. Belongie, J. Malik, and J. Puzicha. Shape matching and object recognition using shape contexts. *Pattern Anal-*

ysis and Machine Intelligence, IEEE Trans., 24(4):509–522, April 2002.
 [3] H. Blum. Biological shape and visual science (part i). *J. Theoretical Biology*, 38(2):205–287, February 1973.
 [4] H. Blum and R. N. Nagel. Shape description using weighted symmetric axis features. *Pattern Recognition*, 10(3):167 – 180, January 1978.
 [5] M. Brady and H. Asada. Smoothed local symmetries and their implementation. *Int. J. Robotics Res.*, 3(3):36–61, September 1984.
 [6] P. Cavanagh. What’s up in top-down processing. In A. Gorea, editor, *Representations of Vision*, pages 295–304. Cambridge University Press, Cambridge, UK, 1991 edition, January 1991.
 [7] T. F. Cootes, C. J. Taylor, D. H. Cooper, and J. Graham. Active shape models - their training and application. *Comput. Vis. Image Underst.*, 61(1):38–59, January 1995.
 [8] A. Dubinskiy and S. Zhu. A multiscale generative model for animate shapes and parts. In *Proc. 9th IEEE ICCV*, volume 1, pages 249–256, October 2003.
 [9] U. Grenander, A. Srivastava, and S. Saini. A pattern-theoretic characterization of biological growth. *Medical Imaging, IEEE Trans.*, 26(5):648–659, May 2007.
 [10] D. D. Hoffman and W. A. Richards. Parts of recognition. *Cognition*, 18(1-3):65–96, December 1984.
 [11] A. Jain, Y. Zhong, and S. Lakshmanan. Object matching using deformable templates. *Pattern Analysis and Machine Intelligence, IEEE Trans.*, 18(3):267–278, March 1996.
 [12] B. B. Kimia, A. R. Tannenbaum, and S. W. Zucker. Shapes shocks and deformations i: the components of two dimensional shape and the reaction diffusion space. *Int. J. Comput. Vision*, 15(3):189–224, July 1995.
 [13] M. Leyton. A process-grammar for shape. *Artificial Intelligence*, 34(2):213 – 247, March 1988.
 [14] S. Mallat and Z. Zhang. Matching pursuits with time frequency dictionaries. *Signal Processing, IEEE Trans.*, 41(12):3397–3415, December 1993.
 [15] F. Mokhtarian and A. Mackworth. Scale-based description and recognition of planar curves and two-dimensional shapes. *Pattern Analysis and Machine Intelligence, IEEE Trans.*, 8(1):34–43, January 1986.
 [16] D. Mumford. Mathematical theories of shape: do they model perception? In B. C. Vemuri, editor, *Geometric Methods in Computer Vision*, volume 1570, pages 2–10. SPIE, July 1991.
 [17] S. Osher and J. A. Sethian. Fronts propagating with curvature-dependent speed. *J. Comput. Phys.*, 79(1):12–49, November 1988.
 [18] T. Pavlidis. *Structural pattern recognition*, volume 1. Springer-Verlag, Berlin, illustrated edition, January 1977.
 [19] E. Sharon and D. Mumford. 2d-shape analysis using conformal mapping. *Computer Vision and Pattern Recognition, IEEE Comp. Soc. Conf.*, 2:350–357, October 2004.
 [20] D. W. Thompson. *On growth and form*. Cambridge University Press, Cambridge, UK, abridged ed./edited edition, January 1961.

- [21] N. Trinh and B. Kimia. A symmetry-based generative model for shape. In *Proc. 11th IEEE ICCV*, pages 1–8, October 2007.
- [22] S.-C. Zhu. Embedding gestalt laws in markov random fields. *Pattern Analysis and Machine Intelligence, IEEE Trans.*, 21(11):1170–1187, November 1999.

Appendix

A. Supplementary Information and Materials

The running time and performance of both shapelet and formlet pursuit algorithms is dependent on multiple parameters as previously described. For the choice of parameters described in Section 5, the 32-component formlet pursuit of an unoccluded shape is computed in approximately 10 minutes running on standard desktop hardware. In comparison, our implementation of shapelet pursuit requires approximately 5 minutes with a similar parameter choice.

We invite the reader to view additional illustrations and movies made publicly available. In addition, our MATLAB implementation of formlets and shapelets, and canonical animal image database used for evaluation, are accessible online. All of these resources and documentation can be found at <http://www.elderlab.yorku.ca/formlets>.

B. Gain Constraint of Diffeomorphic Formlets

Here we address the issue of constraining formlet parameters such that the resulting functions are diffeomorphisms of the complex plane. As the formlets defined in Equation 1 are both isotropic and angle preserving, it is sufficient to require that the radial deformation ρ be a diffeomorphism of \mathbb{R}^+ . This in turn can be enforced by requiring that $\rho(r; \sigma, \alpha)$ be strictly increasing in r . For fixed $\sigma \in \mathbb{R}^+$, we see

$$0 < \frac{\partial}{\partial r} \rho(r; \sigma, \alpha) \quad (10)$$

$$\Leftrightarrow -1 < \alpha \frac{\partial}{\partial r} \sin\left(\frac{2\pi r}{\sigma}\right) \exp\left(\frac{-r^2}{\sigma^2}\right).$$

Finding an explicit solution of this inequality for α is a non-trivial task. However, if we assume $\alpha < 0$, it is easy to see that the minimal slope of ρ is attained when $r \rightarrow 0^+$. Hence, we evaluate Equation 10 at $r = 0$ to find the gain lower-bound:

$$-1 < \alpha \frac{2\pi}{\sigma} \Leftrightarrow -(2\pi)^{-1} < \frac{\alpha}{\sigma}. \quad (11)$$

As we were unable to find a closed-form expression for the upper-bound of α , an approximation was calculated from numerical simulations:

$$-1 < \alpha \frac{\approx -5.1125}{\sigma} \Leftrightarrow \frac{\alpha}{\sigma} \lesssim 0.1956 \quad (12)$$

Together, Equations 11 and 12 provide a necessary and sufficient condition for the monotonicity of the radial deformation ρ . Therefore, by enforcing the constraint of Equation 2 on the gain parameter α , we guarantee that the formlet $f(z, \zeta, \sigma, \alpha)$ is a diffeomorphism of the plane.

C. Computation of Optimal Gain

Suppose that the observed curve Γ^{obs} is approximated by Γ^{k-1} . Given certain formlet scale and space parameters ζ_k and σ_k , we wish to compute the optimal gain

$$\alpha_k = \operatorname{argmin}_{\alpha \in \mathbb{R}} \xi(\Gamma^{obs}, f(\Gamma^{k-1}; \zeta_k, \sigma_k, \alpha)) \quad (13)$$

where, for curves a and b , $\xi(\Gamma^a, \Gamma^b)$ denotes the L_2 error metric

$$\int_0^1 \operatorname{Re}(\Gamma^a(t) - \Gamma^b(t))^2 + \operatorname{Im}(\Gamma^a(t) - \Gamma^b(t))^2 dt. \quad (14)$$

Utilizing the linearized formlet expression of Equation 8, we differentiate ξ with respect to α :

$$\begin{aligned} & \frac{\partial}{\partial \alpha} \xi(\Gamma^{obs}, f(\Gamma^{k-1}; \zeta_k, \sigma_k, \alpha)) \\ &= \frac{\partial}{\partial \alpha} \int_0^1 \operatorname{Re}(\Gamma^{obs}(t) - \Gamma^{k-1}(t) - \alpha g^k(t))^2 \\ & \quad + \operatorname{Im}(\Gamma^{obs}(t) - \Gamma^{k-1}(t) - \alpha g^k(t))^2 dt \\ &= 2 \int_0^1 \alpha (\operatorname{Re} g^k(t))^2 - \operatorname{Re} \Gamma^{res}(t) \operatorname{Re} g^k(t) dt \\ & \quad + 2 \int_0^1 \alpha (\operatorname{Im} g^k(t))^2 - \operatorname{Im} \Gamma^{res}(t) \operatorname{Im} g^k(t) dt \\ & \quad \text{where } g^k(t) = g(\Gamma^{k-1}(t) - \zeta_k; \sigma_k). \end{aligned} \quad (15)$$

Note that Γ^{res} again denotes the residual curve $\Gamma^{obs} - \Gamma^{k-1}$. We then set Equation 15 to zero and solve:

$$\begin{aligned} & \frac{\partial}{\partial \alpha} \xi(\Gamma^{obs}, f(\Gamma^{k-1}; \zeta_k, \sigma_k, \alpha)) = 0 \\ & \Rightarrow \alpha \int_0^1 (\operatorname{Re} g^k(t))^2 + (\operatorname{Im} g^k(t))^2 dt \\ & \quad = \int_0^1 \operatorname{Re} \Gamma^{res}(t) \operatorname{Re} g^k(t) dt \\ & \quad \quad + \int_0^1 \operatorname{Im} \Gamma^{res}(t) \operatorname{Im} g^k(t) dt \\ & \Rightarrow \alpha \|g^k\|_2 = \int_0^1 \langle \Gamma^{res}(t), g^k(t) \rangle dt \\ & \Rightarrow \alpha = \frac{\int_0^1 \langle \Gamma^{res}(t), g^k(t) \rangle dt}{\|g^k\|_2}. \end{aligned} \quad (16)$$

As a result, when ζ_k and σ_k are known, the optimal gain α_k that maximally reduces the L_2 error between the observed curve Γ^{obs} and current approximation Γ^{k-1} is deterministically computed as

$$\alpha_k = \frac{\int_0^1 \langle \Gamma^{obs}(t) - \Gamma^{k-1}(t), g(\Gamma^{k-1}(t) - \zeta_k; \sigma_k) \rangle dt}{\|g(\Gamma^{k-1}(t) - \zeta_k; \sigma_k)\|_2}. \quad (17)$$

Note that with certain space and scale parameters, Equation 17 may produce an optimal gain outside the monotonicity bounds of Equation 2 and Appendix B. When this occurs, the formlet is not guaranteed to preserve topology, and should be ignored. Alternatively, the gain could be set to the nearest constraint endpoint, or some other topology-preserving value.