

Chapter 3

Flux Graphs for 2D Shape Analysis

Morteza Rezanejad and Kaleem Siddiqi

3.1 Introduction

Medial representations, introduced by Blum [2], simultaneously capture properties of an object's outline and its interior. Abstractions of medial representations into graphs have become popular in the computer vision literature and have successfully been applied to view-based object recognition [12, 14]. Recent extensions and applications include alterations of medial graphs to capture salient object parts [8] and the use of medial fragments for perceptual grouping to form object part hypotheses directly from images [11].

Motivated by the success of medial representations, this chapter revisits a quantity related to medial axis computations—the limiting behavior of the average outward flux (AOF) of the gradient of the Euclidean distance function to the object's boundary as the region through which it is computed is shrunk [4]. We exploit the property that at skeletal points the AOF reveals the object angle and thus can be viewed as a scalar descriptor from which the complete boundary can be reconstructed. We then introduce a novel measure of salience for a skeletal point by combining the AOF with a check on uniqueness of the inscribed medial disk to the host skeletal branch. The simplified skeletons are used to derive a directed graph-based representation of the object which we term the flux graph. Our experiments show that flux graphs are a good deal simpler than competing skeletal graphs such as shock graphs, by a number of standard complexity measures, with little loss in representational power. Furthermore, they yield competitive performance in object recognition experiments.

We begin by discussing mathematical properties of the geometry of the medial axis of an object and by introducing the appropriate notation.

M. Rezanejad (✉) · K. Siddiqi
McGill University, Montreal, QC, Canada, H3A 2A7
e-mail: morteza@cim.mcgill.ca

K. Siddiqi
e-mail: siddiqi@cim.mcgill.ca




Fig. 3.1 Local geometry of a maximal inscribed disk centered at the skeletal point \mathbf{p} with radius r and with object angle θ . The maximal inscribed disk touches the boundary at two points $\mathbf{b}^{\pm 1}$ ($\Pi(\mathbf{p}) = \{\mathbf{b}^{+1}, \mathbf{b}^{-1}\}$) (adapted from [13])

Definition 3.1 Assume an n -dimensional object denoted by Ω with its boundary given by $\partial\Omega \in \mathbb{R}^n$. A closed disk $D \in \mathbb{R}^n$ is a *maximal inscribed disk* in Ω if $D \subseteq \Omega$ but for any disk D' such that $D \subset D'$, the relationship $D' \subseteq \Omega$ does not hold.

Definition 3.2 The *Blum medial locus* or *skeleton*, denoted by $Sk(\Omega)$, is the locus of centers of all maximal inscribed disks in $\partial\Omega$.

As illustrated in Fig. 3.1, a skeletal point is characterized by its location \mathbf{p} , the maximal inscribed disk radius r , the object angle θ , the direction of the unit tangent vector \mathbf{T} , and the object angle θ given by $\arccos(-\frac{dr}{ds})$, where s is the arc length along a branch of the medial axis. The projection $\Pi(\mathbf{p})$ is the set of closest points on the boundary $\partial\Omega$ to \mathbf{p} , i.e., $\Pi(\mathbf{p}) \triangleq \{\mathbf{q} \in \partial\Omega : \|\mathbf{p} - \mathbf{q}\| = \min\{\|\mathbf{p} - \mathbf{q}\| \forall \mathbf{q} \in \partial\Omega\}$. For a skeletal point \mathbf{p} the projection set $\Pi(\mathbf{p})$ is the set of points on the boundary touched by the maximal inscribed disk centered at \mathbf{p} (the points $\mathbf{b}^{\pm 1}$ in Fig. 3.1). According to the ‘‘Maxwell set’’ definition of the medial locus [10], each skeletal point $\mathbf{p} \in Sk(\Omega)$ must have at least two closest boundary points ($|\Pi_{\Omega}(\mathbf{p})| \geq 2$).

Topologically $Sk(\Omega)$ consists a set of branches that join to each other at branch points to form the complete skeleton. A *skeletal branch* denoted by χ is a set of contiguous regular points from the skeleton that lie between a pair of junction points, a pair of end points or an end point and a junction point. As shown by Dimitrov et al. in [4] these three classes of points can be analyzed by considering the behavior of the average outward flux of the gradient of the Euclidean distance function to the boundary of a 2D object, given by $\frac{\int_{\partial R} (\hat{\mathbf{q}}, \mathbf{N}) ds}{\int_{\partial R} ds}$, when shrunk to a circular neighborhood, where $\hat{\mathbf{q}} = \nabla \mathbf{D}$ [4], with \mathbf{D} the Euclidean distance function to the object’s boundary. In particular:

1. \mathbf{p} is a *regular point* if the maximal inscribed disk at \mathbf{p} touches the boundary at two corresponding boundary points such that $|\Pi_{\Omega}(\mathbf{p})| = 2$. The computed AOF at a regular point \mathbf{p} is given by $\lim_{\varepsilon \rightarrow 0} \frac{\mathbf{F}_{\varepsilon}(\mathbf{p})}{2\pi\varepsilon} = -\frac{2}{\pi} \sin \theta$.
2. \mathbf{p} is an *end point* if there exists δ ($0 < \delta < r$) such that for any ε ($0 < \varepsilon < \delta$) the circle centered at \mathbf{p} with radius ε intersects $Sk(\Omega)$ just at a single point (r is the radius of the maximal inscribed disk at \mathbf{p}). The computed AOF at an end point \mathbf{p} is given by $\lim_{\varepsilon \rightarrow 0} \frac{\mathbf{F}_{\varepsilon}(\mathbf{p})}{2\pi\varepsilon} = -\frac{1}{\pi} (\sin \theta_P - \theta_P)$.




Fig. 3.2 Different types of skeletal points are illustrated using segments of the skeleton $Sk(\Omega)$ of a given shape Ω . *Left*: A regular skeletal point. *Middle*: An end point. *Right*: A junction point. (Adapted from [4])

3. \mathbf{p} is a *junction point* if $\Pi_{\Omega}(\mathbf{p})$ has three or more corresponding closest boundary points. Generically a junction point has degree 3. All other branch points are unstable. The computed AOF at a junction point \mathbf{p} is given by $\lim_{\varepsilon \rightarrow 0} \frac{\mathbf{F}_{\varepsilon}(\mathbf{p})}{2\pi\varepsilon} = -\frac{1}{\pi} \sum_{i=1}^n \sin \theta_i$.

These different classes of skeletal points are shown in Fig. 3.2.

We now enumerate the main contributions of this chapter. First, previous approaches to compute flux-based skeletons and use them for boundary representation are not entirely complete. Section 3.2 addresses these limitations and presents a method that gives more complete boundary reconstruction results. Second, a new method for skeletal simplification which in turn leads to a simplified graph representation is presented in Sect. 3.3. Underlying this simplification is a measure of saliency that combines a notion of uniqueness of the inscribed medial disk to the host branch with the limiting AOF value.

3.2 Full Boundary Reconstruction

According to the Maxwell set definition of the medial axis, each point on the skeleton has two or more corresponding boundary points. Therefore, given a mapping between boundary points to skeletal points, it is possible to invert that mapping to reconstruct the boundary purely from skeletal points and their properties. Dimitrov et al. [4] attempted to do this by exploiting the relationship between regular points of the medial axis and the object angle. In this section, we will review the basic algorithm for doing this and then extend it to obtain a more complete boundary reconstruction by adding the cases of end points and junction points.

3.2.1 Boundary Representation Through Regular Points with First-Order Approximation of the Tangent Vector

Taking a regular point \mathbf{p} on the skeleton, Dimitrov et al. outlined the reverse transform to obtain corresponding boundary points by $\mathbf{b}^{\pm 1} = \mathbf{p} + r \text{Rot}(\pm\theta) \mathbf{T}_{\mathbf{p}}$. To reconstruct $\mathbf{b}^{\pm 1}$ from a regular point on a parametrized skeleton, the following parameters




Fig. 3.3 *Top row:* Outlines of binary images of a dog, a profile and a hand object, along with their derived skeletons using flux based skeletonization. *Bottom row:* Reconstructed boundary points (filled black disks) overlaid on the original outlines, using the method of Dimitrov et al. [4]

of a skeletal point ought to be numerically computed: the coordinates of the point \mathbf{p} , the radius value r , the object angle θ , and the unit tangent vector \mathbf{T}_p . During the skeletonization process, a parametrized discrete skeleton is computed where each skeletal point includes its position \mathbf{p} , the radius at that point r , and the limiting AOF value. For the object angle θ , a numerical estimate is obtained based on the relationship for regular skeletal points: $\theta = \arcsin\left(-\frac{\mathbf{F}_\varepsilon(P)}{4\varepsilon}\right)$. Finally, the tangent vector is estimated as the slope of the line that connects the prior (discrete) skeletal point \mathbf{p}_{-1} to the subsequent (discrete) skeletal point \mathbf{p}_{+1} , i.e., $\mathbf{T}_p = \frac{\mathbf{p}_{+1} - \mathbf{p}_{-1}}{\|\mathbf{p}_{+1} - \mathbf{p}_{-1}\|}$. Figure 3.3 shows results from these skeletonization and boundary reconstruction algorithms, using the original implementations.

3.2.2 Full Boundary Reconstruction

As is evident from the results in Fig. 3.3, the reconstruction of regular points, though promising, does not provide a complete representation of the boundary. In this subsection, we extend this approach by considering all types of skeletal points and providing a better numerical approximation of the parameters required for reconstruction. To achieve this aim, three limitations of the boundary reconstruction method are considered and addressed:

1. *Sensitivity of first-order approximation of tangent estimation:* The two point stencil computation of the tangent vector is very sensitive to discretization effects along the skeleton, and can often fail at regular points. To mitigate these

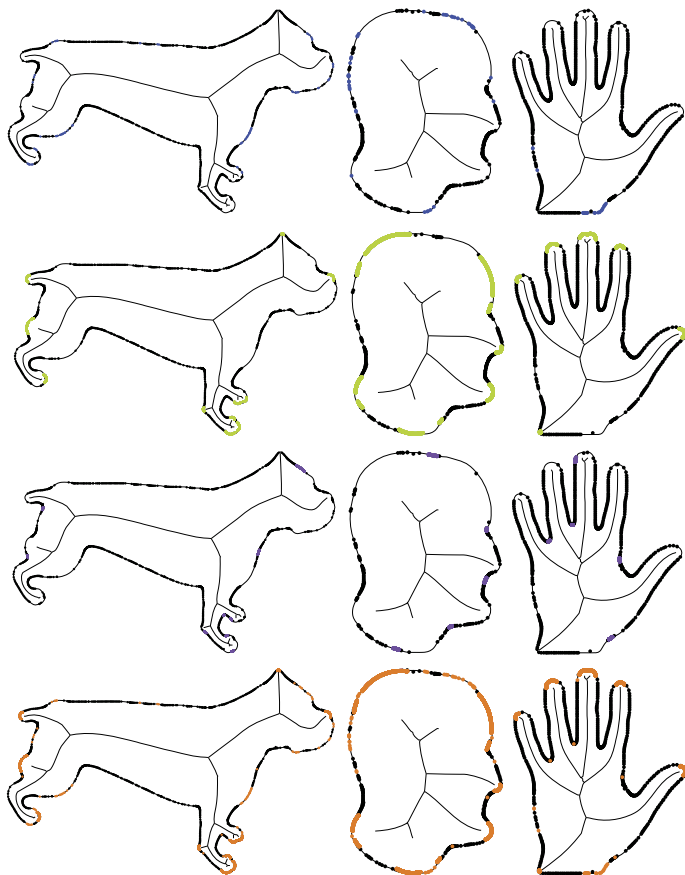


Fig. 3.4 *Top row:* Along with the reconstructed points in Fig. 3.3 shown with black disks, newly reconstructed points resulting from the improved tangent estimation are shown with *blue disks*. *Second row:* Along with the reconstructed points in Fig. 3.3 shown with black disks, newly reconstructed boundary circular segments corresponding to end points are shown with *green disks*. *Third row:* Along with the reconstructed points in Fig. 3.3 shown with black disks, newly reconstructed boundary points corresponding to junction points are shown with *violet disks*. *Bottom row:* Along with reconstructed points in Fig. 3.3 shown with black disks, all the additional reconstructed boundary points are shown in *orange*

numerical errors, we deploy higher order methods for approximating the unit tangent. For those medial loci for which the two point method fails, we use a four point (discrete) stencil approximation [1] given by $\mathbf{T}_p = \frac{2}{3} \left(\frac{\mathbf{p}_{+1} - \mathbf{p}_{-1}}{\|\mathbf{p}_{+1} - \mathbf{p}_{-1}\|} \right) + \frac{1}{3} \left(\frac{P_{+2} - P_{-2}}{\|P_{+2} - P_{-2}\|} \right)$ where \mathbf{p}_{+2} and \mathbf{p}_{-2} represent the subsequent and the previous skeletal points to \mathbf{p}_{+1} and \mathbf{p}_{-1} , respectively. Using the second-order of approximation of tangent estimation results in a number of newly reconstructed boundary points (see Fig. 3.4, top row).

2. *Boundary points that map to an end point*: The boundary reconstruction method by Dimitrov et al. [3] does not explicitly consider the other two types of skeletal points (end points and junction points). This decision results in a number of circular segments missing from the boundary, which map to the end points. We present a numerical approach to recover such missing boundary points. Assume \mathbf{p} is an end point such as the one shown in Fig. 3.2. Then, there would be a circular arc segment from the boundary corresponding to this skeletal point. The osculating disk at \mathbf{p} touches the boundary along that circular segment, and the limiting tangent vector to the skeleton at that point bisects the angle that subtends the circular arc. Let γ represent the curve of that circular arc segment, then

$$\gamma : I \rightarrow \Omega \quad (3.1)$$

$$\gamma(\theta) = \mathbf{p} + r \text{Rot}(\theta) \mathbf{T}_{\mathbf{p}} \quad (3.2)$$

where I is an interval $I = [-\theta_{\mathbf{p}}, \theta_{\mathbf{p}}]$. The coordinates of the point \mathbf{p} , and the radius value r are parameters that are computed during the skeletonization process. To compute γ , the following parameters need to be computed numerically other than \mathbf{p} , and r : the object angle $\theta_{\mathbf{p}}$, and the unit tangent vector $\mathbf{T}_{\mathbf{p}}$. To compute the object angle, we use the end point equation $\frac{F_{\varepsilon}(P)}{2\pi\varepsilon} = -\frac{1}{\pi}(\sin\theta_p - \theta_p)$. For the tangent vector $\mathbf{T}_{\mathbf{p}}$, we simply use the tangent estimation of the (discrete) skeletal point prior to the end point, i.e., $\mathbf{T}_{\mathbf{p}} = \mathbf{T}_{\mathbf{p}_{-1}}$. Figure 3.4 (second row) shows boundary reconstruction results with the newly found circular boundary segments corresponding to end points shown in green.

3. *Boundary points that map to a junction point*: Junction points are also not included in the initial boundary reconstruction method by Dimitrov et al. [3]. We compute the corresponding boundary points of a junction point the same way that we compute the corresponding boundary points of a regular point, with the difference that the tangent vectors near junction points are approximated by those at the prior points on the skeleton. The rest of the procedure is the same as that for computing boundary points for a regular point. Figure 3.4 (third row) shows the improvement with the newly found boundary points corresponding to junction points shown in violet.

The contribution of this approach to reconstructing boundary points is threefold: improved approximation of tangents for many regular points of the skeleton, the computing of circular segments that correspond to end points of the skeleton, and the computing of extra boundary points from junction points. In Fig. 3.4 (bottom row), the additional skeletal points added by these steps are shown in orange, which together with the original reconstructed points demonstrate a far more complete representation of the boundary (compare with Fig. 3.3). The remaining gaps between the reconstructed boundary points can be attributed to the fact that they are mappings of discretely sampled skeletal points.




Fig. 3.5 A part of the dog shape is shown with maximal inscribed disks corresponding to unique and non-unique skeletal points. The maximal inscribed disk centered at $\mathbf{p}_1 (\in \chi_1)$ does not intersect with any maximal inscribed disk from branches other than χ_1 so \mathbf{p}_1 is a unique skeletal point. In contrast, $\mathbf{p}_2 (\in \chi_1)$ is not a unique skeletal point because the maximal inscribed disk centered at $\mathbf{p}_2 (\in \chi_1)$ intersects with the maximal inscribed disk centered at $\mathbf{p}_3 (\in \chi_2)$

3.3 Salient Parts of the Medial Axis

We now build on the previous results to obtain a novel measure of saliency for medial axis points that combines two criteria: (1) The object angle, which by the characterization of [4] is obtained directly from the computation of the AOF and (2) A notion of uniqueness of the maximal inscribed disk at a skeletal point to the host branch.

Definition 3.3 A *unique skeletal point* has the property that the maximal inscribed disk centered at it does not intersect the maximal inscribed disk associated with any skeletal point on any other branch.

Whereas the object angle has often been used as a criterion for saliency [13], the second notion is novel. The intuition here is that unique skeletal points are salient because without them a significant portion of the object's area would not be represented. Examples of unique and non-unique skeletal points are shown in Fig. 3.5.

As explained in Sect. 3.1, the limiting average outward flux at a regular skeletal point \mathbf{p} is computed by: $\lim_{\varepsilon \rightarrow 0} \frac{\mathbf{F}_\varepsilon(\mathbf{p})}{2\pi\varepsilon} = -\frac{2}{\pi} \sin \alpha$. This equation determines a relationship between the AOF and the object angle. The bigger the AOF, the higher the object angle and the more likely the shape silhouette is to be elongated locally. Since elongated parts admit a simple and stable medial axis structure, skeletal points with high AOF are salient.

3.3.1 Simplifying the Skeleton

We combine these two measures of saliency to simplify flux based skeletons using the following procedure: when the considered skeletal point is unique or its normalized AOF is greater than a certain threshold, the skeletal point is retained. In




Fig. 3.6 *Left:* The skeletal points found to be unique are shown in *black* on the medial axis of a dog example. *Middle:* Normalized flux values of a skeleton are shown in a range starting from *white* (minimum AOF) and ending in *black* (maximum AOF). *Right:* Several salient segments labeled as t_i are shown as the result of simplifying the medial axis by retaining only those skeletal points that are unique or have AOF above a threshold

our experiments, we use the threshold $\tau = 0.9045$ for the AOF, which means that all non-unique skeletal points with object angle α greater than about 60° will be retained in the simplified skeleton. Figure 3.6 illustrates the result of applying this simplification procedure on the dog shape.

3.4 Flux Graphs

Our main motivation for simplifying the flux-based skeleton is to extract a graph representation which is simpler than but otherwise as complete and effective as popular existing approaches such as the shock graph [14] and the bone graph [8]. We propose a “Flux Graph” that uses the simplification process to describe a shape as a set of connected parts while preserving the topology of the original skeleton.

3.4.1 Nodes and Edges

The simplification process can result in a number of skeletal fragments, as illustrated by the example in Fig. 3.6. Not all these fragments described distinct parts, rather, those that share a significant portion of their volumes (obtained as the union of the associated medial disks) and are in close proximity of one another can be combined via a merging process. The segments which remain at the end of the merging process are treated as the nodes of a flux graph. The results of merging fragmented parts associated with the simplified skeleton of the dog shape are shown in Fig. 3.7 (left). The set of edges between nodes are then determined based on their connectivities on the original medial axis. To direct edges, we consider the average radii of inscribed disks associated with two adjacent nodes and compare them. The one with larger magnitude is chosen as the parent and the other as the child. The resulting directed flux graph for the dog shape is shown in Fig. 3.7 (right).

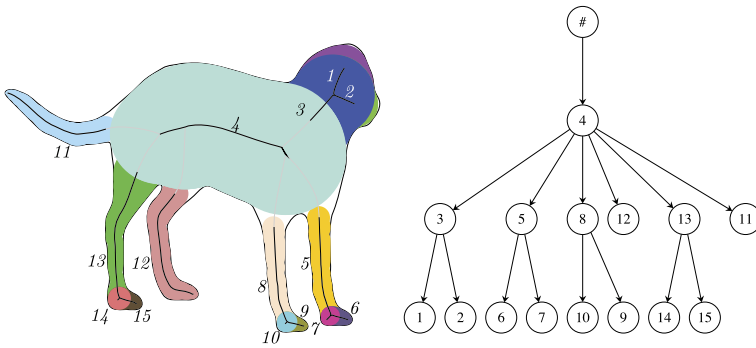


Fig. 3.7 The flux graph of the dog shape. *Left:* The set of nodes is shown with the distinct parts depicted in different colors, each representing a union of medial disks. *Right:* The directed flux graph. The dummy node # carries no geometrical information but serves as a parent to all the top level nodes

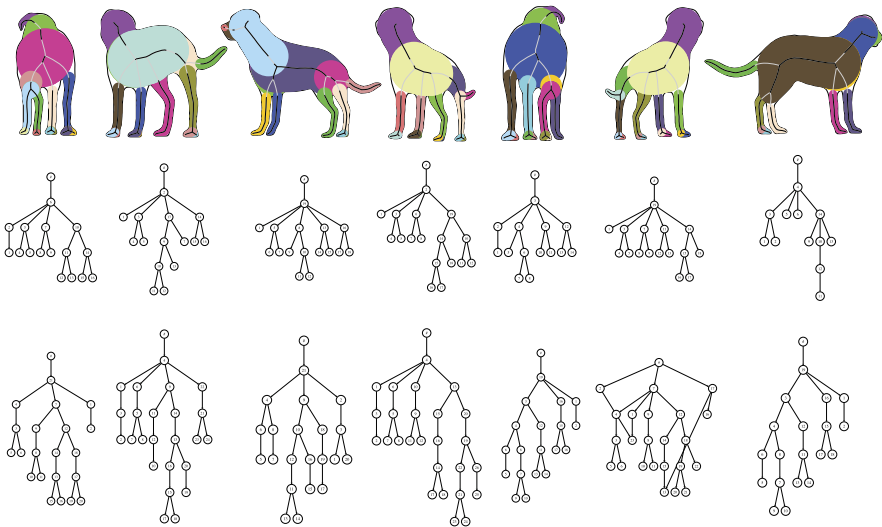


Fig. 3.8 *Top row:* A view of a dog (*middle*) with adjacent views obtained by rotating around it in the clockwise and anti-clockwise directions. For each view, the parts reconstructed by each node of the flux graph are shown as a colored union of disks. *Middle row:* The flux graph corresponding to the view in the *top row*. *Bottom row:* The shock graph corresponding to the view in the *top row*

3.4.2 Qualitative Stability with Viewpoint Changes

We provide a qualitative demonstration that flux graphs remain stable under small changes in viewpoint, while providing an intuitive part structure. We consider a view of the dog (Fig. 3.8 (top row, middle)) and adjacent views obtained by rotating around it in clockwise and anti-clockwise directions. For each view the top row

Table 3.1 Efficiency of flux graphs over shock graphs. The measures in the first six columns are obtained by taking the ratios of the average values of these complexity measures for flux graphs and shock graphs, subtracting these ratios from 1, and then averaging over all the 1664 silhouettes in the database. The last column indicates the percentage of area of the original object reconstructed by flux graphs

	Nodes	Edges	$\Sigma \deg(v)d$	Depth	Skeletal point	TSV	Coverage
Efficiency	% 49.87	% 56.08	% 59.99	% 26.38	% 24.08	% 48.52	% 99

depicts the parts represented by each node of the flux graph, the second row the flux graph and the bottom row the shock graph. Changes to the flux graph typically occur when new parts, such as the tail, come into view (or disappear) but the overall graph structure is much simpler than that of the shock graph. This is essentially because the shock graph utilizes and hence represents the entire skeleton, without any simplification. The experimental results in Table 3.1 which shows averages over 1664 view-based silhouettes of objects used in [8] demonstrate that the flux graph representation is essentially complete, reconstructing 99 % of an object’s area. This will be discussed in further detail in Sect. 3.5.3.

3.5 Flux Graphs for Matching

A skeletal graph abstraction can be used as a tool in many visual shape problems including view-based object recognition. We now examine the potential of using flux graphs for matching, in comparison against the well established shock graphs. To carry out a comparative experiment against shock graphs, we used the same graph matching setup and database used for shock graphs in [5, 14].

3.5.1 Topological and Geometrical Similarity

Given two flux graphs, which are directed acyclic graphs (DAGs) a bipartite graph is constructed between their nodes in a hierarchical manner. Each edge is weighted based on the structural similarity between nodes; the weight is the normalized length of difference of the topological signature vectors (TSVs) introduced in [14]. The best matching of a maximum weighted bipartite matching is when the sum of the values of the edges is maximized. In a DAG representation, the TSV is defined as the vector of eigenvalue-sums derived from the corresponding adjacency matrix for the sub-DAG of the considered node. The matching algorithm used is a greedy algorithm [5] which has the benefit of finding a largest maximal matching in polynomial time. The similarity is computed by matching a query with a model node and then normalizing by the number of matched nodes according to the order of the model graph.

3.5.2 *The DAG Matcher*

To match a query shape with other shapes, we must develop a DAG matcher. The DAG matcher receives two DAGs as input and computes a value representing their similarity, as well as a list of corresponding nodes in the two DAGs. This analysis considers both topological structure (Γ) and geometric information (Δ) associated with a flux graph's vertices. Each of these two measures returns a value normalized in the interval $[0, 1]$. The final similarity score is a weighted combination of these two $S(G_1, G_2) = \omega\Gamma(G_1, G_2) + (1 - \omega)\Delta(G_1, G_2)$, where $S(G_1, G_2)$ represents the similarity between DAGs derived from two given shapes, and ω is a tuning weight in the interval $[0, 1]$. At the end of the process, a list of corresponding nodes and a similarity measure are obtained.

3.5.3 *The Dataset and Experimental Results*


The matching problem we consider is to recognize unseen 2-D query views of 3-D objects by matching a query view against all the available silhouettes (reviewed in Sect. 3.5). We compare results of these experiments with those obtained using shock graphs in [6, 8].

The dataset used for our experiments is the same dataset used for experiments carried out for Bone Graphs in [7, 9] and Shock Graphs [5] and has 13 3-D models. Perspective projection of each 3-D object is computed onto the image plane where each model is centered in a uniformly tessellated view sphere. With 128 uniformly sampled views per object, the data set contains a total of 1664 2-D projected views.

3.5.4 *Flux Graphs versus Shock Graphs*

We begin by demonstrating that by a number of complexity measures the flux graph is simpler and hence more efficient than the shock graph, while essentially providing a complete reconstruction of the original object. To do this, in Table 3.1, for each of the 1664 views we compare: the count of graph vertices, the count of graph edges, the cumulative sum of number of nodes at each depth multiplied by the depth, the depth, total number of skeletal points on the graph, and the average of the TSV (topological signature vectors) values. The numbers reported in the table reflect the efficiency gained by using flux graphs over shock graphs, e.g., flux graphs have 50 % fewer nodes, 56 % fewer edges and 24 % fewer skeletal points. The last column shows the fraction of the area of the original object reconstructed by flux graphs (99 %), indicating that there is essentially no less in representational power.

Fig. 3.9 Using the experimental set up of [6, 8], we compare the use of flux graphs versus shock graphs in a view-based object recognition experiment involving a total of 128 views of each of 13 3-D graphical objects (1664 silhouettes in total). The flux graphs, which are considerably simpler, provide recognition results that are a few percentage points below those of shock graphs



3.5.5 Matching 2-D Views of 3-D Models

We now evaluate the flux graph against the shock graph in a set of view-based object recognition experiments. This comparison follows the matching framework of [7]. The recognition task is performed by: (a) Each view removed sequentially from the database (1664 2-D view-based shapes), and compared to all other remaining views (b) if the class of the closest matching view is the same as that of the query, then the recognition is interpreted as being correct. In the next set of trials, in each step 25 % of the total views are removed randomly from the database. The same experiment is then carried out with further subsampled databases. Figure 3.9 plots the recognition estimation success rates for both shock graphs and flux graphs, averaged over all views of all objects in the database. See [6] for a more detailed explanation of the experimental set up. We also note that the results reported in [6, 8] show that the use of bone graphs, which require a more elaborate construction process, outperforms shock graphs in this experiment.

Flux graphs offer the advantage of efficiency in terms of fewer nodes, edges, depth levels and skeletal points than shock graphs, while still allowing for intuitive hierarchical part-to-part correspondences. However, in terms of the quantitative results, shock graphs outperform flux graphs slightly in this experiment. This could be in part because the matcher used has been tuned to shock graphs and their detailed features and has not been changed in any way to exploit the simplicity of flux graphs. A particular issue is that the geometric node similarity measure used in the matcher [6] implicitly assumes that a node contains a continuous locus of skeletal points. This assumption fails for flux graph nodes that arise from the simplification process we have outlined because the underlying skeletal segments may be fragmented. The greedy matching approach may also suffer from some limitations and alternate hierarchical matching algorithms could be explored.

3.6 Conclusion

We have presented a novel skeletal shape representation that can be used to faithfully reconstruct the original object's boundary from medial entities. The comprehensive recovery of the object's boundary supports the integrity of using the average outward flux at skeletal points for shape analysis. In addition, a complete representation suggests a way of directly relating medial quantities to boundary features, because the medial features are easier to handle, to store and to compare with other represented objects than the shape boundaries directly.

We have suggested the use of the uniqueness of an inscribed disk to the host skeletal branch as a novel measure of saliency. Combining this measure with the limiting AOF leads to simplified skeletons which can be abstracted as graphs that are simpler than popular skeletal graphs in the literature such as shock graphs. In contrast with methods that carry out ligature analysis for simplification based on the limited number of configurations of the placement of ligature and non-ligature parts, such as the bone graph in [8], our investigation has the advantage that the notion of saliency is defined for each skeletal point separately. The flux graph representation has been evaluated using a matching framework designed for shock graphs ([8, 9]) to recognize 2D views of 3D objects and the results show that flux graphs are almost as good as shock graphs for matching. However, more work could be done to improve the robustness of the merging process of fragments left by our simplification method, which is presently based on a heuristic.

To advance the use of flux graphs for matching, a number of directions could be explored including the use of appropriate node similarity measures, the incorporation of a notion of types for nodes (those resulting from simplification, and those not) and the use of alternate hierarchical matching algorithms. The qualitative simplicity and stability of flux graphs with changes in viewpoint suggests their potential for view-based partitioning of the view sphere and view abstraction.

Acknowledgements We are grateful to Diego Macrini for many stimulating discussions, and for sharing his shock graph matching code and his database of 3-D model silhouettes to allow the comparisons in this paper to be carried out. This work was supported by NSERC, the Natural Sciences and Engineering Research Council of Canada.

References

1. Abramowitz M, Stegun IA (1970) Handbook of mathematical functions—with formulas, graphs, and mathematical tables. Dover, New York. 0-486-61272-4
2. Blum H (1967) A transformation for extracting new descriptors of shape. In: Wathen-Dunn W (ed) Models for the perception of speech and visual form. MIT Press, Cambridge, pp 362–380
3. Dimitrov P (2003) Flux invariants for shape. Master's thesis, McGill University, Montreal, July 2003
4. Dimitrov P, Damon JN, Siddiqi K (2003) Flux invariants for shape. In: Proceedings of the IEEE conference on computer vision and pattern recognition, vol 1, pp 835–841
5. Macrini DA (2003) Indexing and matching for view-based 3-d object recognition using shock graphs. Master's thesis, University of Toronto

6. Macrini D (2010) Bone graphs: Medial abstraction for shape parsing and object recognition. PhD thesis, University of Toronto
7. Macrini D, Dickinson S, Fleet D, Siddiqi K (2011) Bone graphs: medial shape parsing and abstraction. *Comput Vis Image Underst* 115(7):1044–1061
8. Macrini D, Dickinson S, Fleet D, Siddiqi K (2011) Object categorization using bone graphs. *Comput Vis Image Underst* 115(8):1187–1206
9. Macrini D, Siddiqi K, Dickinson S (2008) From skeletons to bone graphs: medial abstraction for object recognition. In: *Proceedings of the IEEE conference on computer vision and pattern recognition*, June 2008, pp 1–8
10. Mather JN (1983) Distance from a submanifold in Euclidean space. In: *Proceedings of symposia in pure mathematics*, vol 40, pp 199–216
11. Narayanan M, Kimia B (2012) Bottom-up perceptual organization of images into object part hypotheses. In: *European conference on computer vision*, pp 257–271
12. Sebastian T, Klein P, Kimia B (2004) Recognition of shapes by editing their shock graphs. *IEEE Trans Pattern Anal Mach Intell* 26:551–571
13. Siddiqi K, Pizer SM (2008) *Medial representations: mathematics, algorithms and applications*. Springer, Berlin
14. Siddiqi K, Shokoufandeh A, Dickinson SJ, Zucker SW (1999) Shock graphs and shape matching. *Int J Comput Vis* 35:13–32