

# Appearance Shock Grammar for Fast Medial Axis Extraction from Real Images

Charles-Olivier Dufresne Camaro<sup>1</sup>, Morteza Rezanejad<sup>4</sup>,  
Stavros Tsogkas<sup>1,2\*</sup>, Kaleem Siddiqi<sup>4</sup>, Sven Dickinson<sup>1,2,3\*</sup>

<sup>1</sup>University of Toronto <sup>2</sup>Samsung Toronto AI Research Center

<sup>3</sup>Vector Institute for Artificial Intelligence

<sup>4</sup>School of Computer Science and Centre for Intelligent Machines, McGill University

{camaro, tsogkas, sven}@cs.toronto.edu, {morteza, siddiqi}@cim.mcgill.ca

## Abstract

We combine ideas from shock graph theory with more recent appearance-based methods for medial axis extraction from complex natural scenes, improving upon the present best unsupervised method, in terms of efficiency and performance. We make the following specific contributions: i) we extend the shock graph representation to the domain of real images, by generalizing the shock type definitions using local, appearance-based criteria; ii) we then use the rules of a Shock Grammar to guide our search for medial points, drastically reducing run time when compared to other methods, which exhaustively consider all points in the input image; iii) we remove the need for typical post-processing steps including thinning, non-maximum suppression, and grouping, by adhering to the Shock Grammar rules while deriving the medial axis solution; iv) finally, we raise some fundamental concerns with the evaluation scheme used in previous work and propose a more appropriate alternative for assessing the performance of medial axis extraction from scenes. Our experiments on the BMAX500 and SK-LARGE datasets demonstrate the effectiveness of our approach. We outperform the present state-of-the-art, excelling particularly in the high-precision regime, while running an order of magnitude faster and requiring no post-processing.

## 1. Introduction

Object shape has a fundamental role in visual perception theory. Shape defines a basic level of abstraction that determines the spatial extent of structures in the physical world, and drives object recognition. A popular representation of 2D shape is the *Medial Axis Transform (MAT)* [4].

The medial axis has been of particular interest in both human and computer vision because of its direct relationship to local symmetries of objects. Local symmetries effectively decompose a shape into salient parts, aiding recognition and pose estimation, while being robust to viewpoint changes. At the same time, symmetry in general has been proven to be instrumental in the analysis of complex scenes [26, 32], facilitating the encoding of shape and their discrimination and recall from memory [3, 23, 39]. The importance of symmetry for scene categorization has been recently re-confirmed in [21, 41].

There are many algorithms that compute the MAT of 2D binary shapes. This problem was first discussed by Blum in his seminal work [4, 5], followed by several extensions and variants, including smooth local symmetries [6], shock graphs [25, 31], bone graphs [15, 16, 17] Hamilton-Jacobi skeletons [29], augmented fast-marching [34], hierarchical skeletons [33], and the scale axis transform [10].

In an effort to broaden the application of such methods, interest in the problem of skeleton extraction from natural images has been recently revived, with a focus on using supervised learning. The first such approach is that of Tsogkas and Kokkinos [37], which was later followed by other methods, including the deployment of random forests [35], or convolutional neural networks [8, 11, 13, 14, 27, 40, 43]. Departing from this trend, Tsogkas and Dickinson defined the first *complete* MAT for color images, formulating medial axis extraction as a set cover problem [36]. However, all these recent approaches have an important limitation: medial points are extracted *in isolation*, without explicit consideration of the local context, i.e., the structural constraints imposed by the fact that they must lie on skeletal segments within regions bounded by curves, with the associated generic classification of the medial axis point types [9]. As a result, one has to consider medial proposals at multiple scales for each point, resulting in a very large space of medial point proposals to search. To make things worse,

\*Sven Dickinson and Stavros Tsogkas contributed to this article in their personal capacity as Professor and Adjunct Professor, respectively, at the University of Toronto. The views expressed (or the conclusions reached) are their own and do not necessarily represent the views of Samsung Research America, Inc.

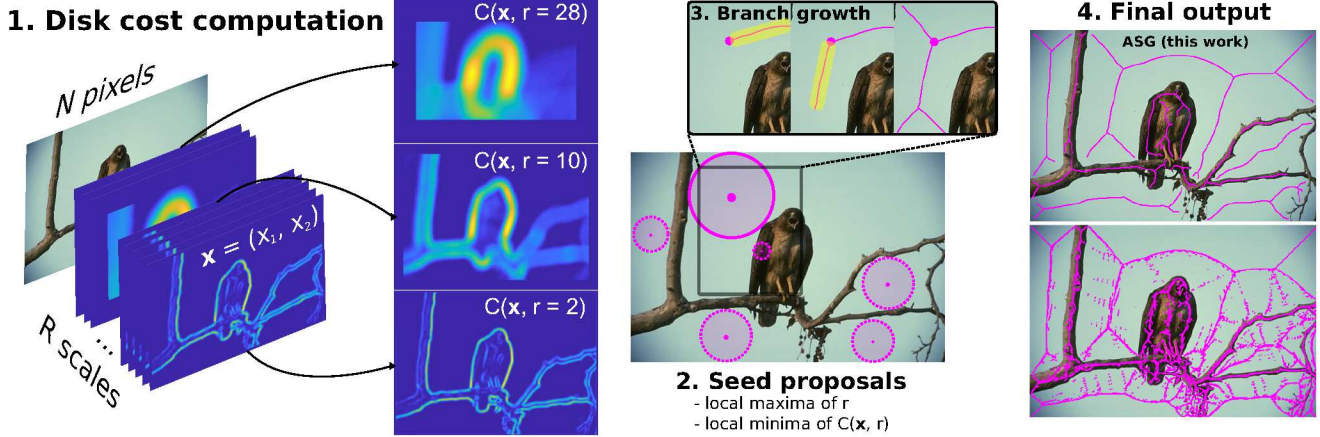


Figure 1: Our ASG algorithm consists of the following steps: **(1) Disk cost computation** associates each valid medial disk proposal with a cost  $C(\mathbf{x}, r)$ . Low costs (blue) represent high “medialness”, whereas high costs (yellow) denote disks that span heterogeneous image regions. **(2) Seed proposals** are selected as local *scale maxima* and local *disk cost minima* (example seeds and the respective disks are shown). **(3) Branch growth** of the selected seed into a medial branch. By following the rules of the SG grammar, the ASG only needs to examine a small, fixed number of proposals in a scale-space neighborhood around a medial point (shown in yellow), making it orders of magnitude faster than the AMAT [36], which naively considers  $O(NR)$  medial disk proposals at each step. **(4) Final output** after growing all seeds. The SG grammar rules automatically enforce connectivity and single-pixel width constraints, producing a sparse, piecewise smooth scene medial axis. In comparison, the AMAT produces much noisier results, that require further post-processing.

post-processing steps, such as non-maximum suppression, are required to ensure 1-pixel width results, or to group medial points into meaningful segments.

In this paper we propose a method that reduces search space redundancies and mitigates the need for post-processing in appearance-based medial point extraction from natural scenes, by combining ideas from shock graph (SG) theory [24, 30, 31] and the AMAT [36]. Specifically, we use the notion of medial disk costs introduced in [36] to come up with new definitions for shock types, tailored to the domain of RGB images. Our new shock type definitions have all the properties of their binary counterparts, unlocking the use of the SG grammar defined in [31]. The grammar allows us to view medial point generation and grouping from natural images as a generative process, in the same spirit as the synthesis of binary shapes via a combination of birth, growth, and death rules, first proposed in [30]. The explicit use of the grammar drastically reduces the size and complexity of the search space and imposes structural regularities in detection, improving performance and computational efficiency. This is visually illustrated in Figure 1.

The benefits of our technique are noteworthy: our *appearance shock grammar* (ASG) results in a  $11\times$  **speed-up** with respect to the original AMAT algorithm, without the need for postprocessing; the use of the ASG ensures that all resulting medial points are connected to form single pixel-wide medial branches. We also raise concerns with the standard evaluation benchmarks that involve multi-

ple scene skeleton ground truths, such as SYMMAX300 [37] and BMAX500 [36]. We propose an alternative evaluation protocol that addresses these issues, and also takes into account the relative importance of individual medial points with regard to boundary reconstruction. On this improved protocol, the ASG outperforms the state-of-the-art in *unsupervised* medial axis extraction by 7.6%, with a sparser and piecewise smoother output.

## 2. Related Work

**Shock-based medial axis extraction.** Blum defined medial axes as the loci of the centers of all disks that can be maximally inscribed in the interior of the shape [4, 5]. An equivalent definition involves *shocks* [12], the points where “grassfire” wavefronts, starting from the boundaries of the shape, meet. Siddiqi *et al.* [30, 31] assume that the shocks composing the medial axis of a bounding contour are first computed and then introduce the concept of *shock graphs*. They “color” shocks into different types according to the local variation of the medial axis radius function, and then define a shock grammar that determines how shocks of different types are connected with one another. The shock grammar can be used to convert a skeleton into a directed acyclic graph, for use by graph matching algorithms to perform shape matching. Shock graphs have also been successfully used in recognition [25] and database indexing [24]. Bone graphs [15, 16, 17] build on shock graphs by decomposing medial axes into object parts, leading to related graphs

for object recognition applications. Medial branches corresponding to salient object parts are tagged as *bones*, while *ligature* segments [2] connect the bones together.

**Medial axis extraction in natural images.** Most recent work on skeleton extraction from natural images relies on supervised learning. Tsogkas and Kokkinos [37] propose a multiple instance learning approach combined with hand-crafted features, tailored specifically to local reflective symmetries. Teo *et al.* [35] improve on this approach by using a more powerful random forest classifier, and by encouraging global symmetric consistency through an MRF representation. Shen *et al.* [27] introduced the first deep-learning approach to solving this problem, where a fully convolutional neural network (CNN) extracts the locations of the skeleton points, while estimating the local medial disk radii, by combining deep features at multiple scales. Ke *et al.* [11] propose a similar framework that stacks Residual Units in its side outputs, improving performance and robustness. In contrast to works that simply fuse (concatenate) side-output responses, Zhao *et al.* [43] create an explicit hierarchy of skeleton features at different scales. This allows for the refinement of responses at finer scales using high-level semantic context, but also of coarser scale responses by using high-detail local responses from early layers of the CNN. Finally, Wang *et al.* [40] frame the skeleton extraction problem as a 2D vector field generation problem using a CNN, where each vector maps an image point to a skeleton point, similar to the Hamilton-Jacobi skeleton algorithm [7, 29].

A completely different, unsupervised approach, the AMAT, was proposed by Tsogkas and Dickinson [36]. The AMAT frames medial axis extraction in color images as a geometric set cover problem and solves it using a greedy approximate solution [38]. The cost assigned to each potential covering element (disk) is provided by a function that prioritizes the selection of maximal disks, leading to a solution approximating the medial axes of structures in the scene.

In the present paper, we use the same concept of disk costs to generalize the definitions of shocks [31] and, in turn, exploit the shock graph theory in the RGB domain. Unlike [31], we do not assume the medial axis is given. Rather, we use the rules of the SG grammar to constrain the number of eligible medial disks that are considered at every step. This allows us to be much more efficient than the AMAT [36], where disks at all possible locations and scales are valid candidates for the greedy algorithm.

### 3. Shock Theory

A shock graph (SG) [30, 31] is a directed acyclic graph (DAG) built from a skeleton. Its nodes correspond to connected components of shocks of the same type, and its edges represent connections between these components. The direction of an edge indicates the direction of the medial axis

radius derivative between the coarser scale and the finer scale shock. The root of the graph is called the *birth shock*.

Shocks represent a colouring for medial points with specific scale (medial axis radius) gradients. A **type 4** shock (blob) corresponds to a single medial point that is a local maximum in scale. Its counterpart, the **type 2** shock (neck), represents a single medial point that is a local minimum in scale and splits its medial branch into separate parts when removed. **Type 3** shocks (ribbons or bends) are sets of connected medial points of equal scales. Finally, **type 1** shocks (protrusion) are sets of connected medial points with monotonically decreasing scales in one direction.

Formally, the shocks can be defined as follows. For a given closed shape  $X$ , let  $M(X)$  be its medial axis representation.  $M(X)$  consists of medial points  $\mathbf{x}$  of scales  $R(\mathbf{x}) \equiv R_{\mathbf{x}}$ . For a medial point  $\mathbf{x} \in M(X)$  and an open disk  $D(\mathbf{x}, \epsilon)$  of radius  $\epsilon$  centered at  $\mathbf{x}$ , let  $N(\mathbf{x}, \epsilon) = M(X) \cap D(\mathbf{x}, \epsilon) \setminus \{\mathbf{x}\}$  represent its  $\epsilon$ -neighbourhood.  $\mathbf{x}$  is

**type 4** if  $\exists \epsilon > 0$  s.t.  $R_{\mathbf{x}} > R_{\mathbf{y}}, \forall \mathbf{y} \in N(\mathbf{x}, \epsilon)$ ;

**type 3** if  $\exists \epsilon > 0$  s.t.  $R_{\mathbf{x}} = R_{\mathbf{y}}, \forall \mathbf{y} \in N(\mathbf{x}, \epsilon) \neq \emptyset$ ;

**type 2** if  $\exists \epsilon > 0$  s.t.  $R_{\mathbf{x}} < R_{\mathbf{y}}, \forall \mathbf{y} \in N(\mathbf{x}, \epsilon) \neq \emptyset$   
and  $N(\mathbf{x}, \epsilon)$  is not connected;

**type 1** otherwise.

While the shock graph represents the relations between connected medial points in terms of their radii, the shock graph grammar reverses the underlying grassfire flow in time. The successive application of its rules defines a generative process that grows parts of an object. The birth rule dictates that *birth shocks can only be types 3 or 4*, while the death rules allow the shock graph to terminate at any shock type. The protrusion rules define how an interval of medial points, with a monotonically changing radius value, can attach at junctions. Finally, the union rules define the conditions under which distinct branches can be connected together.

#### 3.1. Defining Shocks for Natural Images

The ideas presented in Section 3 assume that  $M(X)$  has already been extracted using some skeletonization algorithm. In this work we turn the problem on its head: rather than using the shock grammar to define a graph on a pre-existing skeleton, we use the rules imposed by the grammar to constrain the search space for medial points. To do that, first we have to formally extend the shock type definitions to the domain of natural images. We employ the same notation as in Section 3, introducing new notation when needed.

The key component for determining the coloring of a shock in the binary domain is the computation of  $R(\mathbf{x})$ , the radius of the largest disk, centered at  $\mathbf{x}$ , that remains contained in the open interior  $\mathcal{X}$  of a closed 2D shape. The contour of such disks is tangent to the shape's boundary at



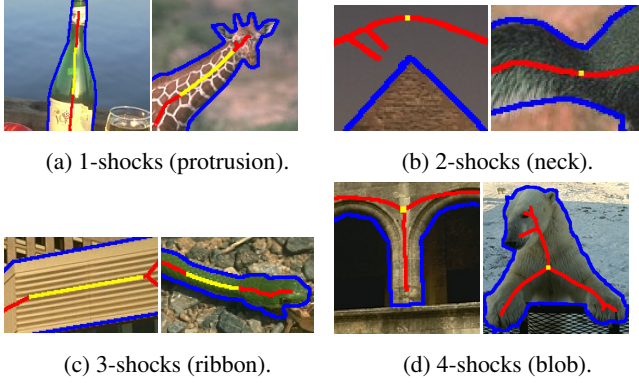


Figure 2: Appearance based shock type examples from BMAX500 [36]. Medial axes are shown in red, contour in blue and selected shocks in yellow.

2 points, at least. Exact computation of  $R(\mathbf{x})$  is feasible because the boundary of a 2D shape is well defined (i.e., the points where the image values change from “0” to “1” or vice-versa). This is not the case in the natural image domain, where extracting object boundaries is an ill-posed problem that typically admits a probabilistic solution.

To deal with this ambiguity, we follow the region-based approach of [36] and assign a cost  $C(\mathbf{x}, r)$  to each disk proposal  $D(\mathbf{x}, r) = D_{\mathbf{x}, r}$ . This cost acts as a “soft maximality” indicator: if  $r$  is close to the ideal (maximal) value,  $C(\mathbf{x}, r)$  is low, whereas disks that are not maximal or cross image boundaries, are severely penalized.

More concretely, let  $\mathbf{x} \in \mathbb{R}^2$ ,  $\mathbf{y} \in N(\mathbf{x}, \epsilon)$  be medial points, and  $R_{\mathbf{x}}, R_{\mathbf{y}} \in \mathbb{R}$  denote the radii of the respective *maximal* disks centered at  $\mathbf{x}$  and  $\mathbf{y}$ . Also, let a small quantity  $\delta_r > 0$  denote an acceptable “cost margin” for determining disk maximality, and  $\epsilon_r > 0$ . Intuitively, if  $C(\mathbf{x}, r + \epsilon_r) - C(\mathbf{x}, r) < \delta_r$ , then  $D_{\mathbf{x}, r + \epsilon_r}$  is a better candidate for being the maximal disk centered at  $\mathbf{x}$  than  $D_{\mathbf{x}, r}$ . We formalize the *scale maximality criterion* as follows:

$$C(\mathbf{x}, R_{\mathbf{x}}) + \delta_r < C(\mathbf{x}, R_{\mathbf{x}} + \epsilon_r). \quad (1)$$

This condition should be satisfied for all disk proposals that are added to our solution. We also define a “cost smoothness” criterion, expressing the fact that the costs of neighboring medial points should not vary significantly. This is another direct analogy to the shock theory for binary shapes, which dictates that the *radii* of neighboring medial points are bound to vary slowly. This is due to the fact that shocks coincide with singularities of a continuous Euclidean distance function from the boundary [5]. Letting  $\delta_c > 0$ , we define the *cost smoothness criterion* as

$$\|C(\mathbf{x}, R_{\mathbf{x}}) - C(\mathbf{y}, R_{\mathbf{y}})\| < \delta_c. \quad (2)$$

By combining these two criteria with the binary shock type definitions, we redefine shock coloring rules in the RGB do-

main. These rules are agnostic to the exact nature of the cost function – we discuss potential choices for  $C$  in Section 4.1. Note, however, that, contrary to the binary case, we must consider all possible locations and scale candidates  $(\mathbf{x}, r_{\mathbf{x}})$ , since we know neither the centers  $\mathbf{x} \in M(X)$  nor the radii  $R$  of the true medial disks. Finally, our shock coloring definitions are adapted to accommodate a discrete pixel grid. For instance, the neighbourhood of a point  $N(\mathbf{x}, 1)$  corresponds to its immediate 8-connected neighbours, while radii only take positive integer values.

## 4. Constrained Medial Point Search Using the Shock Graph Grammar

The formal definition for RGB shocks described in Section 3.1 allows one to use the SG grammar to progressively build an object skeleton while constraining the search space of candidate medial points. We summarize the steps of such an approach in Algorithm 1.

---

### Algorithm 1: Overview of algorithm

---

**Input:** RGB image  $I$

**Output:** Medial points  $M$

- 1 Initialization:  $M \leftarrow \emptyset$
  - 2  $D \leftarrow \text{generateProposals}(I)$ ;
  - 3  $Q_s \leftarrow \text{extractSeeds}(D)$ ;
  - 4 **while** *notEmpty*( $Q_s$ ) **do**
  - 5      $(\mathbf{x}_s, r_s) \leftarrow \text{selectSeed}(Q_s)$ ;
  - 6      $M \leftarrow \text{growSeed}((\mathbf{x}_s, r_s), D)$ ;
  - 7      $Q_s \leftarrow \text{pruneSeeds}(Q_s, M)$ ;
  - 8  $M \leftarrow \text{growEndPoints}(D, M)$ ;
- 

First, we generate medial disk (point) proposals  $D_{\mathbf{x}, r}$  at multiple scales  $r$ . Second, we extract birth seeds  $(\mathbf{x}_s, r_s)$  from the pool of proposals and store them in a queue  $Q_s$ . We grow each seed into a medial axis, by iteratively attaching low-cost medial points<sup>1</sup>. Every time we attach a new point to the axis, we make sure this attachment is consistent with the rules of the SG grammar, and that the medial axis remains connected and one-pixel wide. We greedily continue growing an axis until no points can be added without violating one of these constraints, and then pick the next seed in  $Q_s$  to grow. Note that, because birth seeds can only be type 3 or 4 shocks, which correspond to local scale maxima, the medial axis is constructed in a coarse-to-fine manner. After  $Q_s$  has been exhausted, we relax  $\delta_c$  and grow branch end points that may have been cut short due to the cost constraint. This step allows the algorithm to extend branch growths into more expensive/ambiguous image regions for completeness. We now

<sup>1</sup>In fact, we add medial fragments rather than individual points. See the paragraph on “Seed growth” in this Section.

describe each one of these steps in more detail.

**Proposal generation.** Each medial disk candidate  $D_{\mathbf{x},r}$  is associated with a cost  $C(\mathbf{x}, r)$  that represents how close  $D_{\mathbf{x},r}$  is to being “maximal”. In the domain of real images, a low value for  $C$  is equivalent to a perceptually homogeneous appearance within the disk-shaped region  $D_{\mathbf{x},r}^I \subset I$ . In Section 4.1 we describe in detail two options for  $C$  based on: i) RGB encodings [36]; and ii) image intensity histograms. We compute  $C(\mathbf{x}, r)$  for all points  $\mathbf{x}$  in the image, at all potential scales  $r \in [r_{min}, r_{max}]$ . Proposals corresponding to disks that are not fully enclosed in the image are ignored.

**Seed extraction** should only return type 3 or type 4 shocks. To extract 4-shock *seed* candidates, we scan the space of positions and scales, and check whether the **type 4** criterion holds. For 3-shock *seed* candidates, we check if there is at least one valid neighbour sharing the same scale, as per the shock **type 3** definition. Finally, we impose an additional requirement: a type 3/4 shock  $\mathbf{x}_s$  qualifies as a seed iff it corresponds to a *local cost minimum*, i.e.,

$$C(\mathbf{x}_s, r_{\mathbf{x}_s}) \leq C(\mathbf{y}, r_{\mathbf{y}}), \forall \mathbf{y} \in N(\mathbf{x}_s, 1). \quad (3)$$

All seed candidates are added into a queue  $Q_s$ . Because a 4-seed can eventually grow into a nearby 3-seed as the medial axis is formed (provided that both seeds are part of the same object), once a seed has stopped growing, we also remove any other seeds in  $Q_s$  that have been added to  $M$ .

**Seed selection** follows a coarse-to-fine strategy. We prioritize the selection of seeds with larger radii and lower costs  $C$  because we expect their cost computation to be less sensitive to noise, resulting in more robust axis growth.

**Seed growth** involves attaching medial point proposals to a selected seed  $(\mathbf{x}_s, r_s)$ , following the shock grammar. At each step, the least expensive valid proposal  $(\mathbf{x}, r_{\mathbf{x}})$  in the neighborhood of the axis is added to  $M$ . Proposals whose regions  $D_{\mathbf{x},r}^I$  are subsumed by  $M^I$ , the union of disk regions centered at points in  $M$ , are ignored, as they offer no new information about the object’s shape. The growth process ends when no more valid proposals can be added. To emulate the cost constraint in the RGB shock coloring definitions, we introduce a cost upper bound

$$C_{tol} = C(\mathbf{x}_s, r_s)(1 + \alpha_c) > 0, \quad (4)$$

where  $\alpha_c$  is a small arbitrary positive constant. We ignore proposals with costs larger than  $C_{tol}$  to ensure that the quality of attached points does not degrade during growth.

Single points do not provide sufficient spatial context for determining robust axis growth directions. To resolve these

ambiguities we grow a seed by attaching *fragments* of valid connected medial points,  $F$ , instead. For simplicity, we model medial axis fragments  $F$  as linear segments of length  $l_F \leq l_{max}$ , producing a piecewise-linear approximation of the true medial axis. To rank the quality of candidate fragments we define a *fragment cost*

$$\bar{C}_F = \frac{\alpha(l_F)}{l_F} \sum_{j=1}^{l_F} C(\mathbf{x}_j, r_j), \quad (5)$$

which is proportional to the mean cost of its constituent points. The more expensive a fragment, the less likely it is to be part of the medial axis. To prioritize longer fragments, which provide more context,  $\bar{C}_F$  is weighted by a length-dependent parameter  $\alpha(l_F)$ , i.e., between two fragments with equal mean cost, the longer one will be selected.

At each iteration, we generate multiple candidate fragments and add the one with the lowest  $\bar{C}_F$  to  $M$ . Growth then continues from the endpoint of the last added fragment. This step is repeated until no more valid fragments, i.e., fragments that follow the SG grammar and whose regions are not subsumed by  $M^I$ , can be attached to the current medial branch. In practice this can happen either because the branch is fully grown or because the remaining fragment candidates are too expensive. Then, additional medial branches can be grown from the seed  $(\mathbf{x}_s, r_s)$ .

A medial branch may also terminate at a *junction point*: a medial point from which multiple branches emerge<sup>2</sup>. In this case, new branches can also be grown from that point, as shown in Figure 2d. To identify a junction point, we check if multiple fragments can be attached to it without violating the SG grammar’s protrusion rules.

**End point growth.** Restricting the growth of medial branches using a cost-based threshold for medial fragments promotes robustness and avoids committing to potentially erroneous growth paths. However, the resulting medial axes may not be fully fleshed out: branches corresponding to the fine image details are grown last, and do not always survive this pruning step. To recover these lost medial branches, we perform a final refinement step: we revisit each medial *end point* and allow it to grow further by relaxing the tolerance constraint  $C_{tol}$ , thus allowing less salient fragments to be added. The algorithm terminates when no more valid fragments can be added to any medial end point.

#### 4.1. Cost functions

**Color homogeneity.** We use the default cost function  $C$  of the AMAT [36], after smoothing the input image  $I$  using [42]. The cost of a disk region  $D_{\mathbf{x},r}^I$  with area  $A_r$  is

$$C_{color}(\mathbf{x}, r) = \frac{c(\mathbf{x}, r)}{A_r} + \frac{w_s}{r}, \quad (6)$$

<sup>2</sup>Formally, a junction point is a medial point with at least 3 neighbours.

where  $c(\mathbf{x}, r)$  represents a measure of homogeneity based on  $\mathbf{f}_{\mathbf{x}, r}$ , the average CIELAB space value within  $D_{\mathbf{x}, r}^I$ :

$$c(\mathbf{x}, r) = \sum_k \sum_l \|\mathbf{f}_{\mathbf{x}, r} - \mathbf{f}_{\mathbf{x}_k, r_l}\|_2^2 \forall k, l : D_{\mathbf{x}_k, r_l}^I \subset D_{\mathbf{x}, r}^I. \quad (7)$$

**Intensity histogram.** While straightforward to compute, Equation (6) is sensitive to gradual changes in intensity. We consider a more powerful cost function that is based on local histograms of image intensity and is more appropriate for applications to regions with texture. We first smooth the image using [42]. Then, we precompute a tiling of the image using  $6 \times 6$  squares. For each tile we compute an average intensity value per color channel. We then construct a local histogram  $\mathbf{H}$  for each channel, by placing these averages into one of 10 bins. To compute  $c(\mathbf{x}, r)$ , we replace the  $l^2$ -norm in Equation (7) with a standard Bhattacharyya distance between normalized histograms  $\mathbf{H}_1, \mathbf{H}_2$

$$d_{\text{Bhatt.}}(\mathbf{H}_1, \mathbf{H}_2) = \sqrt{1 - \frac{\sum_i \sqrt{\mathbf{H}_1(i) \cdot \mathbf{H}_2(i)}}{\sqrt{\sum_i \mathbf{H}_1(i)} \cdot \sqrt{\sum_i \mathbf{H}_2(i)}}}, \quad (8)$$

averaged over the 3 color channels, as used for unsupervised texture segmentation in region-based active contours [20]. For each disk under consideration, the histograms are computed using only the enclosed tiles. We also rescale, and add a scale-dependent constant to obtain

$$C_{\text{hist}}(\mathbf{x}, r) = \frac{c(\mathbf{x}, r)}{r} + \frac{w_s}{r}. \quad (9)$$

## 5. Experiments

We conduct experiments on scene and object skeleton detection, on two representative datasets: *BMAX500* [36] and *SK-LARGE* [27]. *BMAX500* is built by automatically extracting skeletons of human-annotated region segments from the *BSDS500* dataset [1]; each image typically comes with 5-7 such annotations. We use the downsampled version of *BMAX500* as in [36], but we also evaluate on the full resolution dataset, to more effectively highlight the computational gains of our approach. *SK-LARGE*, on the other hand, focuses on *object-centric* skeleton detection: each image contains a centered object and the ground truth is only the *foreground* object skeleton. Note that this is a different problem than the one the ASG (and the AMAT) aim to solve, making comparison on the *SK-LARGE* unfair to our algorithm, but we still include it for completeness.

### 5.1. Evaluation Protocol and Criticisms

Traditionally, the evaluation of skeleton extraction methods has followed the protocol originally introduced for the task of boundary detection on the *BSDS500* benchmark [8][18, 19]. According to that protocol, the extracted



Figure 3: Boundary (middle) and skeleton (right) annotations on the *BSDS/BMAX500*. Different colors denote annotations extracted from different segmentations. Whereas boundaries for the same scene form a natural hierarchy, skeletons actually *conflict* with one another, making the evaluation protocol used in [36] unsuitable.




Figure 4: Segmentation (left), binary GT skeletons (middle), and their weighted version based on uniqueness of medial disk area (right) [22]. The most salient skeleton parts are retained (yellow), while skeletal points with low boundary support have low weights (blue).

(boundary/skeleton) map is binarized, and then matched to each one of the available annotations for a given image, using a bipartite graph matching routine that allows for small localization errors. To compute precision (P), a detected point can be matched to *any* of its ground truth (GT) counterparts, while, for perfect recall (R), *all* ground truth points must be matched with a point in the output.

We argue that this benchmarking approach can be misleading for the task of skeleton detection. To better understand why, see Figure 3. The boundary annotations for the same scene form a natural hierarchy: fine-grained interpretations of a scene *complement* the coarser ones, resulting in modest variation in the recall scores. Skeleton annotations, on the other hand, not only change significantly when the source segmentation changes, but actually *conflict* with one another. Even if a predicted skeleton perfectly matches one of the ground truths, it may be at complete odds with the rest, hurting the associated recall and F-score.

Although we employ the same evaluation scheme used in previous work for consistency, we propose the following alternative: for each image, we consider each annotation individually and report scores for the one with the maximum F-score. This is a much more reasonable expectation – we require the output to match *at least one* of the acceptable

Resolution	161 × 241 pixels (half)						321 × 481 pixels (full)					
Method ( $C(\mathbf{x}, r)$ )	AMAT (Color)		ASG (Color)		AMAT (Hist)		ASG (Hist)		AMAT (Hist)		ASG (Hist)	
P	.393	<b>.237</b>	<b>.396</b>	<b>.246</b>	.431	<b>.268</b>	<b>.506</b>	<b>.343</b>	.471	<b>.295</b>	<b>.641</b>	<b>.474</b>
R	<b>.640</b>	<b>.665</b>	.452	<b>.485</b>	<b>.623</b>	<b>.658</b>	.541	<b>.595</b>	<b>.769</b>	<b>.794</b>	.503	<b>.546</b>
F1	<b>.487</b>	<b>.350</b>	.422	<b>.326</b>	.509	<b>.380</b>	<b>.522</b>	<b>.435</b>	<b>.584</b>	<b>.431</b>	.564	<b>.507</b>
*R gains	<b>+0.043</b>	<b>+0.047</b>	+0.032	<b>+0.039</b>	+0.016	<b>+0.020</b>	<b>+0.035</b>	<b>+0.040</b>	+0.018	<b>+0.020</b>	<b>+0.036</b>	<b>+0.040</b>
*F1 gains	+0.012	<b>+0.006</b>	<b>+0.014</b>	<b>+0.008</b>	+0.006	<b>+0.004</b>	<b>+0.016</b>	<b>+0.011</b>	+0.005	<b>+0.003</b>	<b>+0.022</b>	<b>+0.015</b>
t (s)	57.4		<b>7.0 (↓ 8.2×)</b>		33.7		<b>6.5 (↓ 5.2×)</b>		393.2		<b>34.7 (↓ 11.3×)</b>	

Table 1: Results on the BMAX500 using standard evaluation (black) and our proposed [single-annotation protocol](#) (blue). Gains for the ligature-weighted version of BMAX500 are denoted by \*. Timings are averages over the BMAX500 test set. Cost function computation times are excluded from the runtime measurements to compare the two algorithms head to head.

scene interpretations, rather than all of them jointly.

Another observation we make is that large portions of the medial axis may have little to do with boundary reconstruction, but are due to the ligature or the “glue” that holds parts of the object together [5]. Curiously, all studies benchmarked on BMAX500 or SK-LARGE, have ignored this fact. With this in mind, we use a uniqueness of medial disk area-based ligature measure proposed by Rezanejad [22], to weight the contribution of each medial axis point on a scale from  $[0, 1]$ . Figure 4 shows a typical example, where the lower weights near the branch points signal the ligature.

**Parameters** are optimized on the BMAX500 validation set. We use  $\alpha_c = 0.75$ ,  $l_{max} = 10$ .  $\alpha(l_F)$  is set to decrease linearly from  $\alpha(1) = 1$  to  $\alpha(l_{max}) = 0.85$ . We fix these values for all experiments, including those on SK-LARGE.

We use the same values as in [36] for the color cost function, namely  $w_s = 1e - 4$  and the default values for the smoothing operation [42]. For the histogram based cost function, we use  $w_s = 2e - 8$ . Finally, we set  $r \in [2, 41]$  for the half-resolution images and  $r \in [2, 82]$  for the full resolution images. During evaluation, any detected medial point within distance equal to 1% of the image diagonal (in pixels) from the ground truth can be a true positive.

## 5.2. Results

We report quantitative results for scene skeleton extraction, on the half- and full-resolution BMAX500 dataset, in Table 1. We compare the AMAT [36] with post-processing (i.e., grouping and thinning) and the ASG, using the two cost functions described in Section 4.1. We include results for both the standard and our proposed evaluation protocol, as well as the gains due to our ligature weighting.

**The cost function matters.** Using the histogram-based cost increases performance noticeably for both the AMAT and the ASG (+2% and +10% F-score respectively). This result confirms our hypothesis that a powerful cost function that is robust to texture and other local appearance variations is crucial in order to obtain good quality medial axes.

**Performance analysis.** We focus on the results for the intensity-histogram cost function. The standard evaluation protocol rewards the AMAT’s dense yet imprecise output: predicted points have multiple “shots” at matching with one of the multiple GT annotations, and, conversely, a GT point is more likely to match a detected point. This increases recall, making the AMAT perform on par with our method, which produces a much sparser (59% fewer points at full resolution), but *precise* output. Using *one* of the GT per image (blue) calibrates P/R, yielding +5.5% and +7.6% F-score for half and full resolution, respectively, and aligns quantitative results to what we witness qualitatively in Figure 5: a clear advantage of obeying the rules of a shock grammar in skeleton detection. The ASG skeletons are smoother, as singularity theory dictates [9], and less sensitive to boundary artefacts, while maintaining agreement with the ground truth. On the contrary, the AMAT skeletons, where medial hypotheses are evaluated in isolation, contain spurious points and invalid branching topology.

Finally, using a ligature-weighted version of BMAX increases recall for both algorithms, with a net advantage for the ASG, suggesting that branches missed by our method tend to be less important for boundary reconstruction.

**ASG dramatically reduces runtime.** Comparison of the histogram variants of AMAT and ASG in Table 1 show a speedup of 5× for the latter at half-resolution and 11× at full-resolution. Our approach not only is faster by an order of magnitude, it also scales much better with the input image size and the number of scales considered. A detailed breakdown of the algorithm is shown in Table 2.

**Comparison with supervised methods.** In Table 3 we compare to supervised learning methods. SK-LARGE contains annotations only for the foreground object skeletons, so we ignore medial axes outside the object during evaluation. Both the AMAT and ASG produce lower F-scores than Hi-Fi [43] and DeepFlux [40], but this is expected because they are not solving the same problem: the former rely solely on bottom-up features to extract medial axes






Figure 5: Qualitative results. **Left to right:** Ground truth (single annotation), ASG (this work), AMAT [36] (after post-processing). Our method produces sparser, cleaner, and more accurate medial axes, without any post-processing.

Resolution	161 × 241		321 × 481	
Proposal Generation	3.63s	36.0%	63.51s	64.7%
Seed Growth	4.60s	45.6%	18.28s	18.6%
End Point Growth	1.85s	18.3%	15.71s	15.9%
Other	0.01s	0.1%	0.73s	0.8%
Total	10.09s	100%	98.23s	100%

Table 2: Runtime breakdown of the ASG. Timings are averages over the 200 images in the BMAX500 test set. *Other* includes the seed extraction, selection, and pruning steps.

	AMAT [36]	ASG	Hi-Fi [43]	DeepFlux [40]
F1	.509	<b>.511</b>	.724	<b>.732</b>
t (s)	511.9	<b>63.2</b>	0.030 (GPU)	<b>0.019 (GPU)</b>

Table 3: Results on SK-LARGE [27]. Runtimes are averages over the SK-LARGE test images.

of *homogeneous image regions*, whereas the latter incorporate high-level, object-specific information to detect *semantic object skeletons*. Taking these numbers at face value also ignores many “hidden costs” of supervised learning: 1) training deep CNNs for skeleton extraction requires GPUs and segmentations, which are costly and time consuming to collect; 2) these models do not generalize on other datasets: [36] showed that FSDS [28], trained on SK-LARGE, fails to generalize on BMAX500; and 3) they do not easily scale to new classes or granularities; e.g., if a new class is added to the dataset, the model must be retrained.

## 6. Discussion

Our new approach for the efficient extraction of medial axes from cluttered natural scenes uses elements from the shock graph theory of shape. In particular, we have generalized the concept of shocks to the RGB domain by considering region-based cost functions, and have devised an algorithm that leverages the rules of the shock graph grammar to guide the search for medial points. Our approach has several merits: 1) it is fully *unsupervised* and thus can generalize to new datasets without any training; 2) it outperforms the state-of-the-art in unsupervised approaches and is an order of magnitude faster and much more efficient in the number of skeletal pixels generated; and 3) it requires no post-processing, such as thinning or grouping of medial points. In our experiments, we have also raised a concern regarding the way scene skeleton detection frameworks are typically evaluated in our community. To address this, we have proposed an alternative, ligature-based weighted evaluation scheme, that takes into account the relative importance of each medial point for boundary reconstruction, and better reflects performance on benchmarks with multiple ground truth annotations per scene.

## Acknowledgements

We thank the Natural Sciences and Engineering Research Council of Canada (NSERC), the Fonds de recherche du Québec - Nature et technologies (FRQNT), and Samsung for research funding.



## References

- [1] Pablo Arbelaez, Michael Maire, Charless Fowlkes, and Jitendra Malik. Contour detection and hierarchical image segmentation. *IEEE Trans. Pattern Anal. Mach. Intell.*, 33(5):898–916, May 2011. [6](#)
- [2] Jonas August, Kaleem Siddiqi, and Steven W Zucker. Ligature instabilities in the perceptual organization of shape. *Computer Vision and Image Understanding*, 76(3):231–243, 1999. [3](#)
- [3] H.B. Barlow and B.C. Reeves. The versatility and absolute efficiency of detecting mirror symmetry in random dot displays. *Vision research*, 19(7):783–793, 1979. [1](#)
- [4] Harry Blum. A transformation for extracting new descriptions of shape. In *Symposium on Models for the Perception of Speech and Visual Form*, 1967. [1](#), [2](#)
- [5] Harry Blum. Biological shape and visual science (part i). *Journal of theoretical Biology*, 1973. [1](#), [2](#), [4](#), [7](#)
- [6] Michael Brady and Haruo Asada. Smoothed local symmetries and their implementation. *The International Journal of Robotics Research*, 3(3):36–61, 1984. [1](#)
- [7] Pavel Dimitrov, James N Damon, and Kaleem Siddiqi. Flux invariants for shape. In *2003 IEEE Computer Society Conference on Computer Vision and Pattern Recognition, 2003. Proceedings.*, volume 1, pages I–I. IEEE, 2003. [3](#)
- [8] Christopher Funk, Seungkyu Lee, Martin R Oswald, Stavros Tsogkas, Wei Shen, Andrea Cohen, Sven Dickinson, and Yanxi Liu. 2017 ICCV challenge: Detecting symmetry in the wild. In *2017 IEEE International Conference on Computer Vision Workshop (ICCVW)*, pages 1692–1701. IEEE, Oct 2017. [1](#), [6](#)
- [9] Peter J. Giblin and Benjamin B. Kimia. On the local form and transitions of symmetry sets, medial axes, and shocks. *International Journal of Computer Vision*, 54(1-3):143–157, 2003. [1](#), [7](#)
- [10] Joachim Giesen, Balint Miklos, Mark Pauly, and Camille Wormser. The scale axis transform. In *Proceedings of the twenty-fifth annual symposium on Computational geometry*, pages 106–115. ACM, 2009. [1](#)
- [11] Wei Ke, Jie Chen, Jianbin Jiao, Guoying Zhao, and Qixiang Ye. Srn: side-output residual network for object symmetry detection in the wild. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1068–1076, 2017. [1](#), [3](#)
- [12] Peter Lax. Shock waves and entropy. In *Contributions to nonlinear functional analysis*, pages 603–634. Elsevier, 1971. [2](#)
- [13] Chang Liu, Wei Ke, Jianbin Jiao, and Qixiang Ye. Rsrn: Rich side-output residual network for medial axis detection. In *The IEEE International Conference on Computer Vision (ICCV) Workshops*, Oct 2017. [1](#)
- [14] Xiaolong Liu, Pengyuan Lyu, Xiang Bai, and Ming-Ming Cheng. Fusing image and segmentation cues for skeleton extraction in the wild. In *The IEEE International Conference on Computer Vision (ICCV) Workshops*, Oct 2017. [1](#)
- [15] Diego Macrini, Sven Dickinson, David Fleet, and Kaleem Siddiqi. Bone graphs: Medial shape parsing and abstraction. *Computer Vision and Image Understanding*, 115(7):1044–1061, 2011. [1](#), [2](#)
- [16] Diego Macrini, Sven Dickinson, David Fleet, and Kaleem Siddiqi. Object categorization using bone graphs. *Computer Vision and Image Understanding*, 115(8):1187–1206, 2011. [1](#), [2](#)
- [17] Diego Macrini, Kaleem Siddiqi, and Sven Dickinson. From skeletons to bone graphs: Medial abstraction for object recognition. In *2008 IEEE Conference on Computer Vision and Pattern Recognition*, pages 1–8. IEEE, 2008. [1](#), [2](#)
- [18] David Martin, Charless Fowlkes, Doron Tal, and Jitendra Malik. A database of human segmented natural images and its application to evaluating segmentation algorithms and measuring ecological statistics. In *Proceedings Eighth IEEE International Conference on Computer Vision. ICCV 2001*, volume 2, pages 416–423. IEEE, 2001. [6](#)
- [19] David R Martin, Charless C Fowlkes, and Jitendra Malik. Learning to detect natural image boundaries using local brightness, color, and texture cues. *IEEE transactions on pattern analysis and machine intelligence*, 26(5):530–549, 2004. [6](#)
- [20] Oleg Michailovich, Yogesh Rathi, and Allen Tannenbaum. Image segmentation using active contours driven by the bhat-tacharya gradient flow. *IEEE Transactions on Image Processing*, 16(11):2787–2801, 2007. [6](#)
- [21] Morteza Rezanejad, Gabriel Downs, John Wilder, Dirk B. Walther, Allan Jepson, Sven Dickinson, and Kaleem Siddiqi. Scene categorization from contours: Medial axis based salience measures. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2019. [1](#)
- [22] Morteza Rezanejad and Kaleem Siddiqi. View sphere partitioning via flux graphs boosts recognition from sparse views. *Front. ICT*, 2:24, 2015. [6](#), [7](#)
- [23] Fred L Royer. Detection of symmetry. *Journal of Experimental Psychology: Human Perception and Performance*, 7(6):1186, 1981. [1](#)
- [24] Thomas B. Sebastian, Philip N. Klein, and Benjamin B. Kimia. Shock-based indexing into large shape databases. In *European Conference on Computer Vision*, pages 731–746. Springer, 2002. [2](#)
- [25] Thomas B. Sebastian, Philip N. Klein, and Benjamin B. Kimia. Recognition of shapes by editing their shock graphs. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 26(5):550–571, 2004. [1](#), [2](#)
- [26] Daniel Sharvit, Jacky Chan, Hüseyin Tek, and Benjamin B. Kimia. Symmetry-based indexing of image databases. *J. Visual Communication and Image Representation*, 9(4):366–380, 1998. [1](#)
- [27] Wei Shen, Kai Zhao, Yuan Jiang, Yan Wang, Xiang Bai, and Alan Yuille. Deepskeleton: Learning multi-task scale-associated deep side outputs for object skeleton extraction in natural images. *IEEE Transactions on Image Processing*, 26(11):5298–5311, 2017. [1](#), [3](#), [6](#), [8](#)
- [28] Wei Shen, Kai Zhao, Yuan Jiang, Yan Wang, Zhijiang Zhang, and Xiang Bai. Object skeleton extraction in natural images by fusing scale-associated deep side outputs. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 222–230, 2016. [8](#)

- [29] Kaleem Siddiqi, Sylvain Bouix, Allen Tannenbaum, and Steven W Zucker. Hamilton-jacobi skeletons. *International Journal of Computer Vision*, 48(3):215–231, 2002. [1](#), [3](#)
- [30] Kaleem Siddiqi and Benjamin B. Kimia. A shock grammar for recognition. In *1996 Conference on Computer Vision and Pattern Recognition (CVPR '96), June 18-20, 1996 San Francisco, CA, USA*, pages 507–513. IEEE Computer Society, 1996. [2](#), [3](#)
- [31] Kaleem Siddiqi, Ali Shokoufandeh, Sven J Dickinson, and Steven W Zucker. Shock graphs and shape matching. *International Journal of Computer Vision*, 35(1):13–32, 1999. [1](#), [2](#), [3](#)
- [32] Hüseyin Tek, Perry A. Stoll, and Benjamin B. Kimia. Shocks from images: propagation of orientation elements. In *1997 Conference on Computer Vision and Pattern Recognition (CVPR '97), June 17-19, 1997, San Juan, Puerto Rico*, pages 839–845. IEEE Computer Society, 1997. [1](#)
- [33] Alexandru Telea, Cristian Sminchisescu, and Sven Dickinson. Optimal inference for hierarchical skeleton abstraction. In *Proceedings of the 17th International Conference on Pattern Recognition, 2004. ICPR 2004.*, volume 4, pages 19–22. IEEE, 2004. [1](#)
- [34] Alexandru Telea and Jarke J. Van Wijk. An augmented fast marching method for computing skeletons and centerlines. *Eurographics 2002*, 2002. [1](#)
- [35] Ching L Teo, Cornelia Fermuller, and Yiannis Aloimonos. Detection and segmentation of 2d curved reflection symmetric structures. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 1644–1652, 2015. [1](#), [3](#)
- [36] Stavros Tsogkas and Sven Dickinson. Amat: Medial axis transform for natural images. In *Computer Vision (ICCV), 2017 IEEE International Conference on*, pages 2727–2736. IEEE, 2017. [1](#), [2](#), [3](#), [4](#), [5](#), [6](#), [7](#), [8](#)
- [37] Stavros Tsogkas and Iasonas Kokkinos. Learning-based symmetry detection in natural images. In *European Conference on Computer Vision*, pages 41–54. Springer, 2012. [1](#), [2](#), [3](#)
- [38] Vijay V Vazirani. *Approximation algorithms*. Springer Science & Business Media, 2013. [3](#)
- [39] Johan Wagemans. Parallel visual processes in symmetry perception: Normality and pathology. *Documenta ophthalmologica*, 95(3-4):359, 1998. [1](#)
- [40] Yukang Wang, Yongchao Xu, Stavros Tsogkas, Xiang Bai, Sven Dickinson, and Kaleem Siddiqi. Deepflux for skeletons in the wild. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 5287–5296, 2019. [1](#), [3](#), [7](#), [8](#)
- [41] John Wilder, Morteza Rezanejad, Sven Dickinson, Kaleem Siddiqi, Allan Jepson, and Dirk B Walther. Local contour symmetry facilitates scene categorization. *Cognition*, 182:307–317, 2019. [1](#)
- [42] Li Xu, Cewu Lu, Yi Xu, and Jiaya Jia. Image smoothing via l0 gradient minimization. *ACM Trans. Graph.*, 30(6):1–12, Dec. 2011. [5](#), [6](#), [7](#)
- [43] Kai Zhao, Wei Shen, Shanghua Gao, Dandan Li, and Ming-Ming Cheng. Hi-fi: Hierarchical feature integration for skeleton detection. *IJCAI*, 2018. [1](#), [3](#), [7](#), [8](#)