

Medial Axis Isoperimetric Profiles

P. Zhang,¹ D. DeFord,^{1,2} and J. Solomon¹

¹Massachusetts Institute of Technology, USA

²Washington State University, USA

Abstract

Recently proposed as a stable means of evaluating geometric compactness, the isoperimetric profile of a planar domain measures the minimum perimeter needed to inscribe a shape with prescribed area varying from 0 to the area of the domain. While this profile has proven valuable for evaluating properties of geographic partitions, existing algorithms for its computation rely on aggressive approximations and are still computationally expensive. In this paper, we propose a practical means of approximating the isoperimetric profile and show that for domains satisfying a “thick neck” condition, our approximation is exact. For more general domains, we show that our bound is still exact within a conservative regime and is otherwise an upper bound. Our method is based on a traversal of the medial axis which produces efficient and robust results. We compare our technique with the state-of-the-art approximation to the isoperimetric profile on a variety of domains and show significantly tighter bounds than were previously achievable.

1. Introduction

The isoperimetric problem can be traced back to Dido’s problem, in which Dido—the founder of Carthage—wished to maximize her territory while bounding it with a limited supply of strips of bull’s hide [VF83]. This problem is solved by invoking the *isoperimetric inequality*, which states for any shape of perimeter P and area A , $P^2 \geq 4\pi A$; as a result, her encapsulated territory is maximized by a circle.

The isoperimetric inequality motivates the definition of the *isoperimetric ratio* $\frac{4\pi A}{P^2}$, which quantifies compactness of a shape. In the study of geography and political redistricting, the word *compactness* is used informally to distinguish shapes that are not too irregular or distorted, rather than in the formal mathematical sense. For example, a perfect circle is the most compact shape achievable, with isoperimetric ratio 1, while a shape with fractal-esque boundary will have a ratio close to 0. Beyond serving as a basic tool in analytic geometry, the isoperimetric ratio—known as the *Polsby–Popper score* in political science [PP91]—has found application in comparing geographic partitions used in political redistricting.

Applications of the isoperimetric ratio to geographic domains suffer from the “coastline paradox.” In particular, since geographic boundaries can have fractal shapes, their lengths are not well-defined and can change at different length scales. The isoperimetric ratio is unstable to boundary perturbation and will change its compactness score depending on the resolution of the input shape (see Figure 1). This instability can have a severe impact on applications to political redistricting where a low compactness score is frequently presented as evidence of improper intent in the line

drawing process [BNNS19, BS20]. Additionally, when alternative plans are considered in court, scores like Polsby–Popper are used to evaluate whether a proposed plan would be a permissible remedy. While straightforward geometric measures are convenient for making direct pairwise comparisons in court, their simplicity and inherent instability means that they are exploitable and insufficient to distinguish between multiple potential types of undesirable shapes.

Recent works have introduced variations of the classical isoperimetric problem to address the shortcomings above. In particular, the *isoperimetric profile* (IP) modifies the original problem by producing a function rather than a single value. Intuitively, the isoperimetric profile augments the isoperimetric ratio by outputting a geometric compactness score for all length scales rather than at just one. More formally, the profile of a domain $\Omega \subset \mathbb{R}^2$ is defined via the following optimization:

$$\text{IP}_\Omega(t) := \min \{ \text{length}(\partial F) : F \subseteq \Omega \ \& \ \text{area}(F) = t \}, \quad (\text{IP})$$

where F is an inscribed shape within Ω , $\text{length}(\partial F)$ is the perimeter of F , $\text{area}(F)$ is the area of F , and $t \in [0, \text{area}(\Omega)]$ is the multi-scale resolution parameter of the profile. The profile for a given value of t is the smallest perimeter a shape F inscribed in Ω can have while also maintaining area t . There is no restriction that F must be connected. Following the notation in [LS19], we will refer to the F that solves this optimization as $E_\Omega(t)$.

While the isoperimetric ratio is a scalar measurement of compactness, the isoperimetric profile provides an entire plot that measures the *multiscale compactness* of a domain. As such, it is more suited to noisy data such as geographic domains that often exhibit



Figure 1: From left to right: circle, low-res-California, high-res-California, silhouette, perturbed-circle. Their isoperimetric ratios are denoted above them. This indicates that a large range of shapes can lie between essentially two circles on the isoperimetric ratio scale.




Figure 2: Isoperimetric profile plotted for two similar shapes of the same area: (green) the circle, and (red) a perturbed circle. The profile shows that these two shapes are identical at large length scales, but differ significantly on smaller length scales. The blue curve shows what [DLSS19] computes: the convex lower envelope of the profile of the perturbed circle. Note the circles are not plotted on the same scale as the profile.

fractal boundaries on coastlines and rocky territories. The notion of multiscale compactness is demonstrated in Figure 2 for two near-circular shapes that manifest extreme isoperimetric ratio values. While the ratio strongly distinguishes these two shapes, the profile reveals that they are in fact similar at coarse length scales and only differ at fine length scales. Furthermore, the isoperimetric profile is invariant to translation and rotation of the domain. When the areas of different domains are normalized, their isoperimetric profiles can be used as multiscale shape descriptors for measuring similarity and clustering.

Despite the theoretical advantages of the isoperimetric profile, the lack of an efficient algorithm to compute it heavily deters its usage in practice. A recent method takes a convex relaxation of Eq. (IP) using *total variation* (TV) on an Eulerian grid to compute the convex lower envelope of $IP_{\Omega}[t]$ [DLSS19]. This method suffers from several drawbacks: the lower bound is not tight on all valid values of t for any domain; the discretization converges poorly under mesh refinement; and computation of $IP_{\Omega}[t_1]$ does not leverage results from previously computed values $IP_{\Omega}[t < t_1]$, although we note that this final issue could be possibly addressed by using a solution for smaller t as the starting point for the next round of optimization.

The drawbacks in [DLSS19] limit its applicability to accurately

obtaining the isoperimetric profile. To rectify this, we present a novel algorithm that computes upper bounds on the isoperimetric profile that are provably tight for a restricted class of domains. Outside this class, we still attain tight bounds on a subset of scales t . Our algorithm exploits restrictions on the mean curvature of $\partial E_{\Omega}(t)$ [GMT80] as well as recent results on values of $E_{\Omega}[t]$ for special domains [LS19]. This leads us to the construction of our bound via the medial axis of Ω . We use our algorithm to bound isoperimetric profiles for a large variety of domains, including districts, states and countries at various resolutions. We show—on domains for which the isoperimetric profile is known—that our bound is much tighter than bounds computed by previous state-of-the-art.

Our main technical contributions are as follows:

- We generalize previous approaches to understanding the isoperimetric profile by constructing it through the medial axis.
- We design a novel algorithm for efficiently computing upper bounds to the isoperimetric profile.
- We prove that our bound is tight for a restricted set of domains.
- We prove that on general domains Ω our bound is tight for a range of t values derived from properties of Ω .
- We provide experiments demonstrating that our bounds on the isoperimetric profile are tight enough to be used in practice.

2. Related Work

While the original isoperimetric problem has been solved for centuries, many variants remain open and actively studied. For a summary of results on the subject in 2D, on Riemannian 3-manifolds, and on measures, see [Ros05].

A key application of the isoperimetric ratio is as a way to measure geometric compactness to detect gerrymandering [PP91]. In this context, it is known as the *Polsby-Popper* score. In some U.S. states, it is even a legally required statistic to be reported for any proposed districting plan [voMSRP11]. Despite this, [KKKng, DT18, BNNS19, BS20] show that this score is not a reliable measure of compactness. [DT18] suggest a discrete Polsby-Popper score on graphs that alleviates problems with map projection at the cost of geometric sensitivity. [Ehr82] suggest a variation of the isoperimetric ratio that computes the ratio of the area of the full domain to the area of its maximum inscribed circle. Ultimately, this measure ignores high resolution features of the domain. Note that the Ehrenburg test can be extracted from the front of the isoperimetric profile, while the Polsby-Popper score can be extracted from its end with interpolation in between. .

Many variations of the isoperimetric problem differ from the isoperimetric profile we consider in that they use the *relative perimeter* of the inscribed domain rather than *full perimeter*. While the relative perimeter formulations do not count shared boundary between the domain and its inscribed subdomain [Ros05], the full perimeter does. [SZ98] provide conditions under which a perimeter-minimizing subdomain of a convex set is also convex. [LS19] study the maximizers of a curvature functional with close ties to the isoperimetric profile when restricted to domains with no necks. Their results provide a method for computing the isoperimetric profile when restricted to no neck domains. [DLSS19] propose the first algorithmic approach to the isoperimetric profile on general domains by using TV as a generalized measurement of

perimeter of a domain. This relaxation of the problem yields the convex lower envelope of the isoperimetric profile. It remains unknown if computation of the exact isoperimetric profile for 2D polygons is in the polynomial complexity class. Additional open problems for the 2D profile are summarized in [CFG12].

2.1. Properties of solutions to (IP)

While there is no known efficient algorithm to compute the full isoperimetric profile, much is known about $\partial E_\Omega(t)$. In particular, for any fixed value of t , the *free boundary* of $E_\Omega(t)$, $\partial E_\Omega(t) \setminus \partial\Omega$ must have constant and nonsingular mean curvature [GMT80]. In addition, $\partial E_\Omega(t)$ and $\partial\Omega$ must meet tangentially, and any free boundary of $E_\Omega(t)$ must be a circular arc [BES49]. For a small subset of values of t , $E_\Omega(t)$ and $\text{IP}_\Omega(t)$ are known in closed form [LS19]. We will expand on this in §3.3 after establishing more notation and geometric tools.

3. Preliminaries

3.1. Domains

Literature on the isoperimetric profile has typically focused on suitably regular *Jordan domains*.

Definition 3.1 (Jordan Curve [Sul12]). A *Jordan curve* or a *simple closed curve* in the plane is the image of an injective continuous map $\phi : \mathbb{S}^1 \rightarrow \mathbb{R}^2$.

Definition 3.2 (Jordan Domain [Sul12]). A *Jordan domain* is the interior of a Jordan curve.

Beyond restricting attention to Jordan domains, there is typically an additional regularity constraint that $\text{area}(\partial\Omega) = 0$, where $\text{area}(\cdot)$ denotes the Lebesgue 2-dimensional measure. Domains satisfying this condition include any finite polygon and Jordan domains with smooth boundary. For contrast, *Osgood curves* are Jordan curves that can have nonzero area [Kno17], but for our purposes these will not be considered. The set of suitably regular Jordan domains will be denoted \mathbb{J} .

3.2. Necks

We will further refine these domains Ω by the characteristics of their *necks*.

Definition 3.3 (Neck, [LNS17]). A domain $\Omega \subset \mathbb{R}^2$ has a *neck* of radius ρ if there exists a pair of balls $B_\rho(x_0), B_\rho(x_1) \subseteq \Omega$ such that there does not exist a continuous curve $\gamma : [0, 1] \rightarrow \Omega$ satisfying $\gamma(0) = x_0$, $\gamma(1) = x_1$, and $B_\rho(\gamma) \subseteq \Omega$.

Let r_n denote the smallest value of ρ for which Ω has a neck. This definition is built intuitively on the idea of a bottleneck where objects of a certain size are unable to pass through. A domain that does not satisfy Definition 3.3 for any ρ has *no neck*, i.e., $r_n = \infty$. For example, *star* domains have no necks. In contrast, it is clear that the shape in Figure 3 has $r_n < \infty$. This definition does not have a concept of multiple necks nor does it take into account the locations of necks. We will provide a new definition in §6.1 that generalizes necks to incorporate these concepts.

3.3. Morphological Operations

Natural variations of a shape are given by morphological operations. Furthermore $E_\Omega(t)$ for limited values of t can be character-




Figure 3: (Left) The thickly outlined blue shape is the starting domain. The innermost red domain is the erosion of the starting domain by B_r . The outermost green shape is the dilation of the starting domain by B_r . (Middle) Morphological opening of the starting domain by balls of changing radius. Darker blue corresponds to smaller radius. (Right) Morphological closing of the starting domain by balls of changing radius. Darker blue corresponds to smaller radius.

ized by the morphological *opening* of Ω . We will also use morphological *closing* in §7.2 to clean up our results. For these reasons, we provide their definitions; see [Ser00] for more discussion.

Let $B_r(x)$ be a ball of radius r centered at $x \in \mathbb{R}^2$ and $B_r := B_r(0)$. Given a starting domain $\Omega \subset \mathbb{R}^2$, its *erosion* by B_r is

$$\Omega \ominus B_r = \{x \in \Omega : B_r(x) \subseteq \Omega\}. \quad (1)$$

Its *dilation* by B_r is

$$\Omega \oplus B_r = \{x \in \mathbb{R}^2 : \Omega \cap B_r(x) \neq \emptyset\}. \quad (2)$$

Its *opening* by B_r is

$$\Omega \circ B_r = (\Omega \ominus B_r) \oplus B_r. \quad (3)$$

Finally, its *closing* by B_r is

$$\Omega \bullet B_r = (\Omega \oplus B_r) \ominus B_r. \quad (4)$$

These operations and their results are depicted in Figure 3. Intuitively, the opening can be thought of as a blurring of the convex corners of a shape. The closing can be thought of as a way of filling in small holes or smoothing out concave corners of a shape.

With this notation, we recall a property of the isoperimetric profile, namely that $E_\Omega(t) = \Omega \circ B_r$ when $r < r_n$ with corresponding area $t = \text{area}(\Omega \circ B_r)$ [LS19]. It follows that when restricted to no neck domains, the isoperimetric profile can be computed by opening [LS19, Theorem 2.3]. For domains with necks, however, the opening is not generally equal to $E_\Omega(t)$. This is exemplified in Figure 3 where the opening has singular curvature near the neck. Furthermore, we find that in practice the range $t \in [\text{area}(\Omega \circ B_{r_n}), \text{area}(\Omega)]$ accounts for a very small portion of the isoperimetric profile on domains with necks.

3.4. Medial Axis

One of our key ideas is to connect the widely used *medial axis* to the isoperimetric profile. The medial axis allows us to formulate a new definition of neck that will be helpful for our analysis.

Intuitively, the medial axis of a 2D domain is a set of curves that characterize the skeleton of a shape. To formalize this idea, let $\Pi_{\partial\Omega}(x \in \Omega)$ be the projection of x onto the closest boundary point.




Figure 4: (Left) The domain Ω is outlined in black. Connectivity of medial axis edge segments are shown with colors indicating their type. Green edges are edge-edge. Pink edges are vert-vert. Blue edges are edge-vert. This distinction between edge segment types allows us to characterize the shape of the curve (linear or quadratic), as well as its radius function. Brown points indicate medial axis nodes where multiple segments connect. (Right) Piecewise quadratic medial axis is shown with color indicating radius function. Light yellow indicates high radius while magenta indicates low radius.

The medial axis of $\Omega \subset \mathbb{R}^2$ is

$$\text{MA}(\Omega) = \{x \in \Omega \mid \Pi_{\partial\Omega} \text{ is discontinuous at } x\}. \quad (5)$$

Said differently, the medial axis contains all interior points where the closest boundary point is non-unique. The medial axis of a 2D polygon can be decomposed into the union of linear and quadratic edge segments [AVP99]. Each *medial axis edge segment* is a set of points that are equidistant from two boundary objects, *governors*: two boundary vertices (vert-vert), two boundary edges (edge-edge), or one boundary vertex and one boundary edge (edge-vert). These edge segments join at *medial axis nodes*. The *medial axis transform* is the coupling of the medial axis with its radius function $r(x \in \text{MA}(\Omega)) = \text{dist}(x, \Pi_{\partial\Omega}(x))$, which measures its distance to the boundary. The medial axis transform is unique, and it can be used to reconstruct Ω [Blu67]. An example medial axis is provided in Figure 4. Lastly, we denote the the largest circle inscribed in Ω by $\text{inB}(\Omega)$ with corresponding radius $\text{inr}(\Omega) = \max\{r(x) : x \in \text{MA}(\Omega)\}$ and center $\text{inx}(\Omega)$. We assume $\text{inB}(\Omega)$ is unique, which is guaranteed for a generic domain, but not for domains with too much symmetry. In such a case, a randomized boundary perturbation will give us the needed genericity.

4. Motivating the Algorithm

We build our algorithm by first considering ways of computing the isoperimetric profile on domains with no necks, and then extending them to general domains. As mentioned in §3.3, for domains with no necks, $E_\Omega(t)$ is always equal to $\Omega \circ B_r$ for some r . Since $\text{area}(\Omega \circ B_r)$ decreases strictly monotonically with r , r and t are in correspondence. Thus, we can build the isoperimetric profile by measuring the area and perimeter of $\Omega \circ B_r$ for values of $r \in [0, \text{inr}(\Omega)]$. It remains to consider different algorithmic ways of building $\Omega \circ B_r$. A key consideration of our algorithm will be how well it extends to general domains with necks.

4.1. Direct computation of $\Omega \circ B_r$

The opening of a 2D domain is simple to compute from modern geometry processing tools and generates feasible solutions to prob-

lem (IP) on general domains. Unfortunately, the IP upper bound generated by measuring the perimeter of $\Omega \circ B_r$ exhibits several undesirable traits on domains with necks. Firstly, $\Omega \circ B_r$ varies discontinuously with r , leaving large gaps in the profile; we repair this issue in §7.2. Second, this IP bound is not monotonic, and we know the true isoperimetric profile should always monotonically increase. Finally, $\partial(\Omega \circ B_r)$ can have singular curvature on free boundaries, while $\partial E_\Omega(t)$ cannot. For these reasons direct computation of $\Omega \circ B_r$ does not extend well to domains with necks. To fix these shortcomings, we propose a novel medial axis based approach.

4.2. Medial Axis Limited Reconstruction

Recall from §3.4 that the medial axis transform is invertible i.e. we can reconstruct Ω from $\text{MA}(\Omega)$ and its radius function $r(x)$. Unsurprisingly, performing the reconstruction on a subset of $\text{MA}(\Omega)$ will produce a subset of Ω . We will refer to these as *limited reconstructions*.

Definition 4.1 (Limited Reconstruction). Let $g \subseteq \text{MA}(\Omega)$ be a subset of the medial axis. The medial axis *limited reconstruction* of Ω by g is $\Omega_g = \bigcup_{x \in g} B_{r(x)}(x)$.

When $g = \text{MA}(\Omega)$, $\Omega_g = \Omega$. It turns out that for a particular choice of g , we can relate the morphological opening of a domain to its limited reconstruction.

Proposition 1. For a radius parameter $\rho \in [0, \text{inr}(\Omega)]$ and $g_\rho = \{x \in \text{MA}(\Omega) \mid r(x) \geq \rho\}$,

$$\Omega_{g_\rho} = \Omega \circ B_\rho. \quad (6)$$

For proof, see §A. The relationship between morphological operations and the morphological skeleton has been studied in image processing [Lan77], where a similar statement is made in the discrete setting. Recall that the isoperimetric profile of no neck domains in \mathbb{J} can be constructed by morphological opening. By Proposition 1 these can now be constructed from subsets g of the medial axis. Furthermore, g can be built iteratively due to the following result:

Proposition 2. For generic Ω and $\rho_1 \geq \rho_2$, we have $g_{\rho_1} \subseteq g_{\rho_2}$. As extreme cases, $g_0 = \text{MA}(\Omega)$ and $g_{\text{inr}(\Omega)} = \text{inx}(\Omega)$.

Proof. If a point $x_1 \in \mathbb{R}^2$ is in g_{ρ_1} , it must be on the medial axis and have $r(x_1) \geq \rho_1$. Then we also have $r(x_1) \geq \rho_2$, so it must be in g_{ρ_2} . At the smallest valid value $\rho = 0$, g_0 is the entire medial axis, because the radius function is non-negative. At the largest nontrivial value $\rho = \text{inr}(\Omega)$, $g_{\text{inr}(\Omega)}$ only contains points at least $\text{inr}(\Omega)$ away from the boundary. For generic Ω , this is $\text{inx}(\Omega)$ by definition. \square

This motivates our algorithm in §5 for constructing the isoperimetric profile by starting with $g_{\text{inr}(\Omega)}$ and iteratively increasing g by greedily absorbing neighboring medial axis points of highest radius. We will see in §7 why this medial axis based approach far outperforms direct computation of $\Omega \circ B_r$.

4.3. Differential change in isoperimetric profile when traversing the medial axis

Since our algorithm will traverse the medial axis, we consider here how much the isoperimetric profile of Ω_g will change as g grows

differentially. We can simplify this problem by looking at only a single medial axis edge segment $\gamma: [0, 1] \rightarrow \mathbb{R}^2$, and let $g(\tau_2) = \bigcup_{\tau \in [\tau_1, \tau_2]} \gamma(\tau)$, for $\tau_1, \tau_2 \in [0, 1]$ and $\tau_1 < \tau_2$. We derive that

$$\lim_{\tau_2 \rightarrow \tau_1} \frac{\text{length}(\partial\Omega_{g(\tau_2)}) - \text{length}(\partial\Omega_{g(\tau_1)})}{\text{area}(\Omega_{g(\tau_2)}) - \text{area}(\Omega_{g(\tau_1)})} = \frac{1}{r(\gamma(\tau_1))}. \quad (7)$$

For derivation see “differentialGrowth.nb” in supplemental materials. Equation (7) tells us that if we infinitesimally increase g by absorbing vertices of the medial axis near x , the slope of the isoperimetric profile will be $1/r(x)$. A consequence of Equation (7) is that the differential growth of the isoperimetric profile when growing a circle of fixed center from radius 0 is infinite, i.e., differentially, it is costly for F in problem (IP) to be disconnected.

5. Algorithm

Our algorithm, as motivated by §4.2, can be roughly thought of as a greedy traversal of the medial axis. Given a domain Ω , we initialize $g = \{\text{inx}(\Omega)\}$ and allow it to iteratively absorb adjacent points of the medial axis of maximal radius. As $|g|$ grows, $\text{area}(\Omega_g)$ and $\text{length}(\partial\Omega_g)$ form our IP upper bound. To put this into practice, we must first discretize several quantities.

First, we assume the input domain Ω is a polygon provided in the form of a $n \times 2$ vertex matrix $\bar{\Omega}$ (not a closure). Next we discretize the medial axis $\text{MA}(\bar{\Omega})$ into a graph MAG whose nodes densely sample the exact medial axis (details on this sampling below). Thus, g is simply a subgraph of MAG. Lastly, we discretize the limited reconstructions F as polygons $\bar{F} = \bar{\Omega}_g$. This is computed as

$$\bar{\Omega}_g = \bigcup_{e \in g. \text{edges}} \bar{\Omega}_e, \quad (8)$$

where each $\bar{\Omega}_e$ is the limited reconstruction from one edge of the medial axis graph. $\partial\bar{\Omega}_e$ consists of two semi-circular arcs and the linear subsets of $\partial\bar{\Omega}$ that are governors of e . Therefore $\partial\bar{\Omega}_g$ is also a union of linear and circular arcs. We linearize circular arcs by polylines such that the length of any edge is less than input parameter d_x . As $d_x \rightarrow 0$, \bar{F} approaches F .

The sampling of $\text{MA}(\bar{\Omega})$ into a graph MAG is performed with respect to two parameters d_c and r_l . The *radius limit* r_l is the lower threshold radius for when we truncate the medial axis. In order to make sure we do not lose any information about necks of the domain, we always set $r_l \leq r_n$. In our experiments, $r_l = r_n$, unless we are on a no neck domain, in which case $r_l = .1$. For any choice of $r_l \leq r_n$, our algorithm can bound the profile for $t \in [0, \text{area}(\Omega \circ B_{r_l})]$. Our discretization of the medial axis also discretizes our upper bound on the IP into a piecewise linear fit. If e is a medial axis graph edge with endpoints x_1 and x_2 , based on equation (7) we enforce that the difference between $\frac{1}{r(x_1)}$ and $\frac{1}{r(x_2)}$ is less than d_c . This concentrates samples of the IP upper bound to regions with high second derivative. As d_c decreases, the resolution of our piecewise linear IP upper bound increases.

Thus our algorithm is as follows. We first compute the exact medial axis of Ω using CGAL’s edge Voronoi procedures [Kar20]. Next, the medial axis is discretized into a graph MAG as described



Figure 5: Sequence of inscribed subdomains $F \subset \Omega$ for (IP) visualized as a colormap. From left to right the domains shown are **California, Boxes, District 1, Worm, and Mozambique**. Each point is colored by the earliest time on the IP bound when F contained that point. The top row shows results from morphological opening (with modification from §7.2), while the bottom row shows results from Algorithm 1.

above. We then initialize $g^0 = \{\text{inx}(\Omega) \in \text{MAG}\}$ and greedily absorb neighboring graph nodes and edges of maximal radius per iteration. In each iteration, \bar{F}_i is the limited reconstruction $\bar{\Omega}_{g^i}$. Our algorithm outputs the sequence of \bar{F}_i , their areas t_i and their perimeters p_i . This procedure is summarized by Algorithm 1. Figure 5 provides a visualization of the \bar{F}_i constructed by this procedure.

5.1. Direct $\Omega \circ B_r$

For comparison with the direct computation of $\Omega \circ B_r$ suggested in §4.1, we pick 100 evenly spaced values of $r_i \in [0, \text{inr}(\Omega)]$, compute $\Omega \circ B_{r_i}$ and its corresponding area and perimeter. This is easily implemented with Matlab’s `polybuffer` methods and forms a second upper bound that in the majority of test cases performs worse than our medial axis based approach. In very few cases, it can perform better, in which case our final IP upper bound is the minimum of these two upper bounds.

6. Theoretical Tools

Our algorithm in the previous section produces a sequence of inscribed shapes \bar{F} . We would like to analyze how close they are to optimal. To do that, we first establish some theoretical tools that will help us analyze the behavior of our algorithm. These tools bridge the gap between the definition of necks and the medial axis.

6.1. Medial Axis Necks

Since our algorithm is derived jointly from the medial axis and the concept of necks, we create a generalized definition for necks based on the medial axis transform. From this definition, we can quantify the size of multiple necks as well as their locations (see Figure 6). This will allow us to prove results that are not possible to state with Definition 3.3.

Definition 6.1 (Generalized Necks). $\eta(\Omega)$ is the set of $x \in \text{MA}(\Omega)$ that satisfy either of the following:

1. x is on the interior of a medial axis edge segment, and $r(x)$ is a local minimum at x .
2. x is on a medial axis node and there are two distinct outgoing medial axis edge segments e_1, e_2 on which the change in $r(x)$ in both of those directions is positive.

Algorithm 1 Isoperimetric Profile Upper Bounder

```

1: procedure COMPUTE-IP-BOUND( $\bar{\Omega}$ ,  $d_c$ ,  $r_l$ )
2:    $[\text{MA}(\bar{\Omega}), r(\cdot)] \leftarrow \text{GETMEDIALAXIS}(\bar{\Omega})$ 
3:    $\text{MAG} \leftarrow \text{MATOGRAPH}(\text{MA}(\bar{\Omega}), d_c, r_l)$ 
4:    $g^0 \leftarrow \text{inx}(\bar{\Omega})$ 
5:    $\bar{F}_0 \leftarrow \text{inB}(\bar{\Omega})$ 
6:    $t_0 \leftarrow \text{area}(\bar{F}_0)$ 
7:    $p_0 \leftarrow \text{length}(\partial\bar{F}_0)$ 
8:   while  $g^i \neq \text{MAG do}$ 
9:      $v \leftarrow \text{GETADJACENTNODES}(\text{MAG}, g^i.\text{nodes})$ 
10:     $c \leftarrow v \setminus g^i.\text{nodes}$ 
11:     $c^* \leftarrow \text{ARGMAX}(r(c))$ 
12:     $e^* \leftarrow \text{GETEDGE BETWEEN}(c^*, g^i)$ 
13:     $g^{i+1} \leftarrow \text{ADDNODE}(g^i, c^*)$ 
14:     $g^{i+1} \leftarrow \text{ADDEDGE}(g^{i+1}, e^*)$ 
15:     $\bar{F}_{i+1} \leftarrow \bar{\Omega}_{g^{i+1}}$ 
16:     $t_{i+1} \leftarrow \text{area}(\bar{F}_{i+1})$ 
17:     $p_{i+1} \leftarrow \text{length}(\partial\bar{F}_{i+1})$ 
18:  end while
19:  return  $t, p, \bar{F}$ 
20: end procedure

```




Figure 6: Domains with medial axes indicated by colored curves and necks depicted by isolated black points. (Left) Domain containing 2 necks of different sizes (Right) Pinched annulus with one neck by definition 6.1, but no necks by definition 3.3

The size of a neck $x \in \eta(\Omega)$ is $r(x)$.

Algorithmically, the computation of the necks $\eta(\Omega)$ is done by first computing the medial axis, and then iterating over its edges and nodes while checking the neck conditions. We denote the maximum neck radius by $r_m = \max\{r(x) : x \in \eta(\Omega)\}$. Similarly we denote its minimum neck radius by $r_n = \min\{r(x) : x \in \eta(\Omega)\}$. If a domain has no necks then we say $r_n = \infty = r_m$. The maximum neck radius r_m is not part of the previous definition of necks but will be used in Proposition 4 to give theoretical guarantees about our algorithm. A nice property of our neck definition is that it is compatible with the previous neck definition, i.e., the minimal neck radius of both are identical.

Proposition 3 (Equality of necks). For $\Omega \in \mathbb{J}$, r_n of Definition 3.3 is equivalent to r_n of Definition 6.1.

For proof see §B. Note that this equivalence does not hold for domains outside \mathbb{J} . Consider the annulus with a pinch in it, illustrated in Figure 6. This does not have a neck by Definition 3.3, but by Definition 6.1 it has a single neck marked by the black point.

6.2. Thick Necks

While much is known about the isoperimetric profile for shapes with no necks ($r_n = \infty$), much less is known when r_n is finite. One might hope that if r_n were in some sense *large enough*, we could still compute the isoperimetric profile with some confidence.

Here we define the *thick neck* condition which will be used in Proposition 5 to analyze the behavior of our algorithm on domains whose necks are large enough. Recall that $\text{inB}(\Omega)$ is the largest circle inscribed in Ω .

Definition 6.2 (Second Largest Circle). Let the radius of the second largest circle in Ω without replacement be

$$\text{inr}_2(\Omega) = \text{inr}(\Omega - \text{inB}(\Omega)). \quad (9)$$

Definition 6.3 (Thick Neck). A neck $x \in \eta(\Omega)$ is *thick* if

$$r(x) > \frac{\text{inr}_2(\Omega)}{2} \quad (10)$$

Definition 6.4 (Thick Neck Condition). A domain satisfies the *thick neck* condition if all of its necks are thick.

For examples of domains that satisfy the thick neck condition and their profiles, see Figure 12.

7. Analysis of Algorithm 1

Given the tools of the previous section, we can now analyze how the perimeter of our constructed shapes \bar{F}_i compares to that of $\partial E_\Omega(t)$. At the t_i 's where our algorithm samples \bar{F}_i , and p_i , our bound satisfies several properties.

7.1. Theoretical analysis

By construction, Algorithm 1 produces an upper bound to the isoperimetric profile. This is because \bar{F}_i is always a feasible candidate to (IP) with $t_i = \text{area}(\bar{F}_i)$. By applying Equation (7) to domains in \mathbb{J} , our bound is strictly monotonically increasing, similarly to the actual IP. Our construction also satisfies the condition in [GMT80], which states that $\partial E_\Omega(t)$ must lie tangent to $\partial\Omega$ and its free boundary must have no points of singular curvature. This is because the free boundary of \bar{F}_i is made of circular arcs that start and end on $\partial\Omega$ and whose centers are on different edges of the medial axis. If these circular arcs intersect to form a point of singular curvature, then a non-contractable loop can be drawn in Ω by tracing from the intersection through both edges of the medial axis, back to $\text{inx}(\Omega)$, implying $\Omega \notin \mathbb{J}$.

We provide guarantees for when our bound is tight for any domain in \mathbb{J} . While we cannot guarantee that our bound is tight for all t_i , by Proposition 4, we know our bound is tight for t_i in a limited range.

Proposition 4. Let Γ_r be the connected component of $\Omega \ominus B_r$ that contains $\text{inx}(\Omega)$. Let $r_t = \max(r_m, \text{inr}_2(\Omega)/2)$. For any domain in \mathbb{J} , our bound for $IP_\Omega(t_i)$ is tight for $t_i > \text{area}(\Omega \circ B_{r_n})$ and $t_i < \text{area}(\Gamma_{r_t} \oplus B_{r_t})$.

Proof. For $t_i > \text{area}(\Omega \circ B_{r_n})$, we know that $E_\Omega(t_i) = \Omega \circ B_{r_n}$. In this region, the remaining area in Ω not covered by $E_\Omega(t_i)$ has no more necks. Thus the problem reduces to the case with no necks, on which our algorithm produces tight bounds. For proof of the latter range of t_i values, see Appendix C. \square




Figure 7: Bounds on the IP computed by our medial axis traversal algorithm on the **Candy** domain. By Proposition 4 our bound is tight on the green regions (left of “Conservatively tight” and right of “Minimal Neck”). On the rest of the profile, our medial axis traversal algorithm is an upper bound. Visualizations of $F_i = \Omega_g$ are depicted along the profile bound.

While Proposition 4 does not cover the entire IP, it applies to any domain in \mathbb{J} regardless of the presence or size of necks. To achieve this result, this proposition uses r_m , a quantity that is not defined in prior work.

Application of Proposition 4 is demonstrated in Figure 7. We compute our upper bound to the IP by Algorithm 1 on the **Candy** domain. Its profile is depicted between two events: the area of the “max inscribed circle” area ($\text{inB}(\Omega)$), and the “max area” area (Ω). We apply Proposition 4, which tells us that our medial axis bound for the IP is tight to the left of the “Conservatively tight” purple event line and to the right of the “Minimal Neck” purple event line. For values of t_i between the two purple lines, we upper-bound the IP.

Given some assumptions on the domain Ω , we show that our IP bound exactly recovers the IP.

Proposition 5. If $E_\Omega(t)$ grows only continuously except when transitioning to being disconnected, then for $\Omega \in \mathbb{J}$ that satisfy the thick neck condition, our bound is tight for all t_i .

Intuitively, this is because the smallest perimeter cost for having a disconnected $E_\Omega(t)$ is larger than the cost for continuing to grow $E_\Omega(t)$ via limited reconstructions. For proof, see §D. Proposition 5 allows us to compute the exact isoperimetric profile for domains satisfying the thick neck condition. We show in §8.3 examples of domains satisfying this condition and that while previous methods of computing the isoperimetric profile are not tight, our method provides the exact profile for this larger class of domains.

7.2. Empirical Quality Improvements

Following Equation (8) we compute Ω_g by the union of floating precision polygons. This sometimes introduces artifacts like holes, or jagged boundary at a scale that is invisible. While seemingly negligible, many holes or jagged regions add up to significantly increase the perimeter thus loosening our upper bound. This is solved easily by some basic post processing. First we remove all holes in F , then we perform morphological closing with a ball of small radius, and lastly we intersect it with Ω to make it feasible again.




Figure 8: (Left) \bar{F}_i before post processing. Several holes and boundary flaws are visible by magnification. (Right) \bar{F}_i with post processing. Holes and boundary flaws are repaired. Post processing changes the area by less than $1e-5$, while decreasing the perimeter by 0.1.




Figure 9: (Left) Our post processing procedure gets rid of small floating point artifacts in the inscribed shape resulting in lower perimeter and a tighter IP bound. (Right) The profile generated by $\Omega \circ B_r$ discontinuously jumps from the green event line to the red event line. $\Omega \ominus B_r$ is visualized in red and shows the creation of a very small disconnected region at x^- . We interpolate the profile in between these discontinuous transitions by inflating a circle at x^- from radius 0 to r .

These steps are summarized in Algorithm 2. Figure 8 shows the change in \bar{F}_i before and after post processing. Figure 9 demonstrates the difference in the IP with and without post processing.

We also perform a repair step to the upper bound obtained by direct computation of $\Omega \circ B_r$. Since $\Omega \circ B_r$ changes discontinuously with r for domains with necks, this leaves gaps in the IP. An example of this discontinuous transition is visualized in Figure 9. The discontinuous change happens due to the creation of a new disconnected point x^- in $\Omega \ominus B_r$. $\Omega \circ B_r$ will instantly include a circle centered on x^- of radius r , which abruptly increases area and perimeter. We detect when the number of disconnected regions in $\Omega \ominus B_r$ changes, and compute x^- as the center of the bounding box of the smallest area component of $\Omega \ominus B_r$. We then inflate a disk centered at x^- from radius 0 to r . This procedure continuously fills in the IP bound and is illustrated in Figure 9.

Algorithm 2 PostProcess Polygon

```

1: procedure POSTPROCESSPOLYGONS( $F, r_c, \Omega$ )
2:    $F \leftarrow \text{REMOVEHOLES}(F)$ 
3:    $F \leftarrow F \bullet B_{r_c}$ 
4:    $F \leftarrow F \cap \Omega$ 
5:   return  $F$ 
6: end procedure
    
```




Figure 10: We test three different discretizations of total variation (TV). (1) TV-1: the TV discretization of [DLSS19]. (2) 3x3 Gradient: TV where gradients are computed with a 3x3 stencil. (3) 5x5 Gradient: TV where gradients are computed with a 5x5 stencil. We test how well these three discretization perform at approximating the perimeter of a polygon when applied to an indicator function of it on finer grid resolutions. While none of the three converge exactly to the correct perimeter, the 3x3 gradient performs significantly better.

8. Experiments and Results

8.1. Comparison Methods

We compare our algorithm against the few existing computational approaches to the isoperimetric profile. First we compare Algorithm 1 with direct morphological opening. While both generate upper bounds, direct opening has much fewer guarantees and generally produces a looser bound. Next we compare with the TV approach [DLSS19] which generates lower bounds. In theory, it generates the convex lower envelope of the IP.

Code for our algorithm, figure generation, and all parameters used can be found in supplementary materials. We show profiles starting from the area of the maximum inscribed circle because the profile before then is uninformative.

8.1.1. Modification to comparison method

For comparisons with TV [DLSS19] we use a modification of their code that significantly improves their lower bound. Their algorithm hinges on the use of total variation (TV) on indicator functions as a measurement of perimeter. Their discretized TV comes from [CDPP16], which we find converges poorly to perimeter. Instead, we implement a discretization based on smoothed second




Figure 11: We compute three different isoperimetric profile bounds on domains with no neck: **Star**, **no-neck-J**. Both morphological opening and our medial axis algorithms produce tight upper bounds to the IP. TV roughly captures the convex lower envelope of the profile but slightly overshoots due to discretization error.

order finite differences:

$$\frac{\partial}{\partial x} = \frac{1}{6 \cdot dx} \begin{bmatrix} -1 & 0 & 1 \\ -1 & 0 & 1 \\ -1 & 0 & 1 \end{bmatrix}, \quad \frac{\partial}{\partial y} = \frac{1}{6 \cdot dy} \begin{bmatrix} 1 & 1 & 1 \\ 0 & 0 & 0 \\ -1 & -1 & -1 \end{bmatrix} \quad (11)$$

The total variation is then $\sum_{i,j=1}^n \left\| \frac{\partial u}{\partial x}, \frac{\partial u}{\partial y} \right\|_2 dx dy$. A comparison of the different discretizations at perimeter estimation are shown in Figure 10. We also tested a smoothed third order finite difference gradient but find that the second order gradient achieves the closest perimeter estimate by far. We use this modified TV for all comparisons with [DLSS19]. We sometimes obtain results where the lower bound intersects the upper bound. This is because even our modified TV does not exactly converge pointwise to continuous TV. Design of such a discretization is the subject of active research and is a challenge with TV-based algorithms more broadly.

8.2. No Neck Domains

On no neck domains both the morphological opening procedure and our medial axis traversal will produce the exact isoperimetric profile. This is depicted in Figure 11. We test on a star domain **Star** as well as a non-star, no neck domain **no-neck-J**. We choose $r_l = .1$ which computes most of, but not the entire, profile. This is why our profile stops before the maximum area is reached. The total variation lower bound roughly follows the convex lower envelope




Figure 12: We compute three different isoperimetric profile bounds on domains with thick necks: **Jellyfish**, **Hourglass**, and **Worm**. Morphological opening computes loose upper bounds, and TV computes loose lower bounds. While morphological opening can produce the exact IP on the small interval between “Minimal Neck” and “Max area”, only our medial axis based algorithm computes the exact IP for the entire range of areas.

of our profile, but overshoots in a few regions. As a lower bound, it leaves slack room, while our algorithm is exact.

Note that the choice of $r_l = .1$ is specific to the domains we tested on. As r_l roughly corresponds to the smallest length scale of a domain our IP bound will capture, a smaller domain requires a smaller r_l .

8.3. Thick Neck Domains

On thick neck domains our algorithm computes the exact IP, while morphological opening computes a much looser upper bound, and total variation computes the convex lower envelope of the profile. This is demonstrated in Figure 12 on the **Jellyfish**, **Hourglass**, and **Worm** domains. Note that the “Minimal Neck” area, $\text{area}(\Omega \circ r_n)$, is often very close to the maximum area. Prior methods would only be able to compute the exact profile in the limited region from $\text{area}(\Omega \circ r_n)$ to $\text{area}(\Omega)$, while our method computes the entire profile. Due to the presence of necks, direct morphological opening produces much worse upper bounds than our method.

8.4. General Domains

We test on a set of geographic and hand drawn data. Maps of **Alabama**, **California**, **Delaware**, and **District 1** of Alabama were obtained from [FIP20]. A map of **Mozambique** was obtained from [Vem20]. In addition, we create a drawing, **Boxes**, to test on. **Alabama**, **California**, and **Delaware** satisfied the thick neck condition while **District 1**, **Mozambique**, and **Boxes** do not. Our profiles for these domains are shown in Figure 13. Our bound computes the exact IP for all t on **Alabama**, **California**, and **Delaware**. [DLSS19] continues to provide convex lower envelopes with slack and an occasional intersection with the upper bound due to imperfect discretization. This lower bound is particularly loose, on **Alabama**, and **California** revealing a significant gap for most of the profile. Morphological opening produces upper bounds as well, but are loose compared to our medial axis based approach and has no theoretical guarantees except for a small region on the tail of the profile.

On the **District 1**, **Mozambique**, and **Boxes** domains, we use Proposition 4 to compute “Conservatively tight” and “Minimal Neck” events that allow us to guarantee our bound is tight on a larger range of t values. For both **District 1**, and **Mozambique**, our algorithm produces much tighter upper bounds than morphological opening. An exception occurs on **Boxes**, where the morphological opening produces a tighter bound. This is explained by the fact that the optimal $E_\Omega(t)$ is disconnected for much of the profile, while our algorithm by construction only produces connected polygons. For this reason it would be worthwhile to explore using disconnected medial axis reconstructions based on a criteria like the thick neck condition. Further exploration of disconnected limited reconstructions is discussed in §9 and left to future work.

8.5. Aggregate Profile

By combining the bounds and guarantees provided in this paper and those of [DLSS19, LS19], we compute the envelope inside which the isoperimetric profile must lie. This is visualized in Figure 14. We find that this produces a relatively narrow region of uncertainty

Number of Vertices	Mesh name	MA time (s)	Open time (s)	TV time (s)	r_n
9	no-neck-J	29.99	6.185	26.6	∞
11	Star	11.489	1.44	13.9	∞
14	Candy	18.694	7.72	115.9	0.444
43	Delaware	303.52	7.69	24.4	∞
46	Jellyfish	156.81	5.70	165.3	0.069
146	Mozambique	1251.1	18.50	10.6	0.040

Table 1: Runtimes of our method, and comparison methods in seconds. Since our method is Lagrangian using floating precision geometry, its runtime is correlated with number of vertices. Total variation uses a pixelization of the domain, and we choose 100 pixels vertically for all experiments listed. Morphological opening is also Lagrangian and shows slight correlation with number of vertices. Our runtime is comparable to the runtime of TV on moderately sized polygons.

where the IP is unknown. This visualization also makes clear the rich space of multi scale behaviors that different domains can have.

For example, consider the **Jellyfish** (brown). It consists of a large cap region which has mostly low resolution features, followed by several tendrils that require high resolution data to capture. This is reflected in its isoperimetric profile where it starts shallow for most of the area, before sharply veering up to have one of the largest perimeters in this figure. **Delaware** (bright green) demonstrates similar behavior to a lesser extent. It starts with a shallow profile indicating its coarse lower portion, before its profile curves upwards to capture the high resolution details of its upper quarter. **Delaware** requires less high resolution data to capture than the **Jellyfish**, however, as indicated by its profile sitting below that of the **Jellyfish**. Our results with the isoperimetric profile make it possible to computationally and numerically compare these domains' multi-scale properties.

8.6. Runtime

A list of runtimes for our experiments is provided in Table 1. The runtime of our algorithm depends most strongly on number of vertices $|V|$, and loosely on other parameters d_c , r_l , and r_n . For morphological opening, the runtime depends primarily on the number of samples of the IP that are computed. For all times listed, we sample 100 values of the IP. For TV, runtime depends on the number of samples of the IP that are computed as well as the pixel resolution used to rasterize the polygon. We compute 20 samples of the TV IP and use 100 vertical pixels for the rasterization.

9. Discussion and Conclusion

This work enables computation of a significant portion of the isoperimetric profile, a far more robust measurement of geometric compactness than the popularly used isoperimetric ratio. We stress that while more challenging to compute, a multi-scale approach to isoperimetry is necessary to perform reliable shape analysis. By providing an upper bound counterpart to [DLSS19] we hope to lower the barrier between users and multi-scale measurement of geometric compactness.

This complementary bound and the relatively small gap between

them in many of our examples as in Figure 14 provides significantly increased confidence that the IP can be used in practice for analyzing potential districting plans. While the examples in [DLSS19] demonstrated that the IP could reveal multiscale structure in individual districts and highlight distinct types of undesirable behavior in district shapes, the fact that the TV relaxation only recovers the lower convex envelope means there was no guarantee that the computed function was representative of the actual IP.

The quite tight IP bounds of **District 1** in Figure 13 give an example of several features common among real districting plans, including the small size of the maximum inscribed circle and rapid acceleration of the IP for large areas as the fine details are filled in. In addition to the problems of coastline and resolution discussed in the introduction, districting plans are usually constructed out of discrete units like census blocks and hence the map drawer may have little control over the finest level of details, which are most heavily penalized by scores like Polsby-Popper that report only a single number to summarize the compactness of the plan. In between the extremal regions, the profile recovers the relative size of necks and tendrils, which are in some cases more likely markers of gerrymandering, for example by packing together urban centers, than rough boundary elements.

While we have made several advancements on this problem, there remain many avenues in which to approach further improvement. On the theoretical side, we conjecture that $E_\Omega(t)$ can always be constructed from a union of balls centered on the medial axis. None of our many experiments have demonstrated otherwise. Proof of this conjecture would decisively reduce the complexity of this problem from 2D to just 1D.

Empirically, our medial axis based algorithm suffers when the optimal solution has disconnected $E_\Omega(t)$. Many heuristics are available for constructing disconnected medial axis reconstructions and merit further exploration. More recent work in computer graphics have used semidefinite relaxations of quadratic problems with great success. This could be formulated into a method of selecting a subset of all balls centered on the medial axis for the reconstruction. A more brute force approach could be to compute the next largest inscribed circle at each iteration and compare that with connected traversal of the medial axis. Lastly, our definition of necks partitions Ω into a union of no neck subregions. This suggests a coarser scale optimization where the atomic unit is a no neck subregion rather than just a ball.

On the lower bound side, [DLSS19] produces a convex relaxation of the problem where mass density can take values between 0 and 1. A simple way to tighten this lower bound is just to run a subsequent nonlinear optimization initialized at the relaxed solution. This would be reminiscent of techniques used in topology optimization for mass constrained mechanical part design.

Our medial axis based approach can be naively extended to 3D, where it also produces upper bounds. Unfortunately in the 3D setting $\partial E_\Omega(t)$ is not necessarily a spherical shell but rather a minimal surface. It would be interesting to measure how big the gap between upper and lower bounds are in the 3D setting.

The isoperimetric profile can be used to distinguish between shapes that are intrinsically identical. For example, [GW96, Figure




Figure 13: Isoperimetric profile bounds computed on a mix of sketch and map data. (Left column) *Alabama, California, and Delaware.* These U.S. states satisfy the thick neck condition so our upper bound is tight, while other bounds are significantly looser. (Right column) *Mozambique, District 1, and Boxes.* These domains do not satisfy the thick neck condition. Our bound still provides a tighter upper bound than opening on *Mozambique, and District 1,* but is looser on *Boxes.* This looseness is explained by the optimal inscribed shape being disconnected, something our algorithm can not achieve by construction. By Proposition 4, our bound is tight to the left of “Conservatively tight” and to the right of “Minimal Neck”.




Figure 14: We combine bounds and guarantees of this paper and those of [DLSS19, LS19] to produce a tight regime inside which the IP must lie. Domains are scaled to have area 1 for plotting purposes. We find that this produces a relatively narrow range of uncertainty that illuminates the rough shape of the IP. This plot also allows us to see the rich space of multi-scale phenomena that 2D domains can have. In particular, the (brown) *Jellyfish*'s profile starts shallow before sharply veering upwards, reflecting the composition of the *Jellyfish* by a big round cap and many small tentacles.

15] define different shapes D_1 and D_2 that have identical spectra. However, their isoperimetric profiles are guaranteed to be distinct since $\text{inr}(D_1) \neq \text{inr}(D_2)$. This simple demonstration incentivizes further exploration of the isoperimetric profile as a shape descriptor.

By relating three classic geometry processing concepts, the medial axis, morphological opening, and isoperimetry, we achieve significant progress towards computation of the isoperimetric profile.

Acknowledgements

The authors thank Edward Chien, Giorgio Saracco, and Hugo Lavenant for many valuable discussions. Paul Zhang acknowledges the generous support of the Department of Energy Computer Science Graduate Fellowship and the Mathworks Fellowship. Daryl DeFord and Justin Solomon acknowledge the generous support of the Prof. Amar G. Bose Research Grant. Justin Solomon and the MIT Geometric Data Processing group acknowledge the generous support of Army Research Office grants W911NF1710068 and W911NF2010168, of Air Force Office of Scientific Research award FA9550-19-1-031, of National Science Foundation grant IIS-1838071, from the MIT-IBM Watson AI Laboratory, from the Toyota-CSAIL Joint Research Center, from a gift from Adobe Systems, and from the Skoltech-MIT Next Generation Program.

Appendix A: Proof of Proposition 1, $\Omega_{g_\rho} = \Omega \circ B_\rho$

Proof. Recall that by definition $g_\rho = \{x \in \text{MA}(\Omega) : r(x) \geq \rho\}$ is the set of medial axis points whose radii are more than ρ . An alternative description for $\Omega \circ B_\rho$ is the union of all circles of radius ρ that fit in Ω . It suffices to show that Ω_{g_ρ} is also the union of all circles of radius ρ that fit in Ω .

Given a point $x \in \Omega_{g_\rho}$, it must be contained in a ball of radius larger than or equal to ρ in Ω because Ω_{g_ρ} is the union of such circles. Such a circle must be in $\Omega \circ B_\rho$ and so x must be as well.

Start with a point $x \notin \Omega_{g_\rho}$. Consider the biggest circle in Ω that contains x . If the radius of this circle is larger than or equal to ρ , then this circle can be enlarged until it lies tangent to multiple boundary vertices while still containing x . The center of the enlarged circle would be on the medial axis of radius more than ρ , which contradicts $x \notin \Omega_{g_\rho}$. Therefore there is no circle of radius larger than or equal to ρ that encloses x i.e. $x \notin \Omega \circ B_\rho$. \square

Appendix B: Proof of Proposition 3: Neck Equality

Proof. We start with the point $x_0 \in \text{MA}(\Omega)$ of minimal neck radius. We can assume x_0 satisfies condition 2 in Definition 6.1 because condition 1 can be reframed into condition 2 by adding an imaginary medial axis junction node at x_0 . Label the two outgoing medial axis edge segments on which change in $r(x)$ is positive by e_1 , and e_2 .

First we show that Ω has a neck of radius $r(x_0)$ by Definition 3.3. This is constructed by taking x_1, x_2 on e_1, e_2 respectively an ϵ distance away from x_0 . By Definition 6.1, $r(x_1) > r(x_0)$, and $r(x_2) > r(x_0)$. Since Ω is a Jordan domain, the path from x_1 to x_2 that stays furthest from $\partial\Omega$ must pass through the choke point x_0 . Even then, at the choke point, a ball of radius $r(x_1)$ or $r(x_2)$ will not fit since $r(x_0)$ is smaller. Therefore, Ω contains a neck of radius $r(x_0)$.

Next we show that Ω does not have a neck of radius $\rho < r(x_0)$. Given any x_1, x_2 and balls of radius $\rho < r(x_0)$ around them, we can construct γ satisfying Definition 3.3 by the following. Split γ into three parts. Let the first part be γ_1 the shortest path contained in Ω from x_1 to the medial axis. Denote the point where γ_1 contacts the medial axis by x_1^m . Let the third part be the reverse of the shortest path γ_3 contained in Ω from x_2 to the medial axis. Denote the point where γ_3 contacts the medial axis by x_2^m . Finally let the second part γ_2 be the path from x_1^m to x_2^m via the medial axis. Since γ connects the two balls, Ω does not contain a neck of radius ρ . \square

Appendix C: Proof of Proposition 4

Proof. Starting from $\text{inB}(\Omega)$, we can increase $E_\Omega(t)$ continuously until the influence of the first neck is reached. This corresponds to the first time g absorbs a point on the medial axis of radius r_m . At this event, $\Omega_g = \Gamma_{r_m} \oplus B_{r_m}$. Before then, $E_\Omega(t)$ is expanding in a neckless domain and so, excluding the possibility of generating a new disconnected region, greedy traversal is optimal.

Taking into account the possibility of generating a new region, the earliest moment could happen is if the radius of curvature of $\partial E_\Omega(t)$ is equal to $\frac{\text{inr}(\Omega)}{2}$. Therefore, to be conservative, let $r_t = \max(r_m, \frac{\text{inr}(\Omega)}{2})$. Until $t = \text{area}(\Gamma_{r_t} \oplus B_{r_t})$ it will be optimal to proceed as if in a no neck domain. \square

Appendix D: Proof of Proposition 5: Tight upper bound for thick neck domains assuming $E_\Omega(t)$ grows only continuously excluding disconnected circles

Proof. We consider ways in which $E_\Omega(t)$ can increase.

A new disconnected component could be generated anywhere in $\Omega - E_\Omega(t)$. In such a case, by the isoperimetric inequality, the optimal shape of the new component is a circle of maximal radius. For an optimistic measurement of how preferable this option is we replace $\Omega - E_\Omega(t)$ by $\Omega - \text{inB}(\Omega)$. The radius of the next largest circle that could be inscribed in $\Omega - \text{inB}(\Omega)$ is by definition $\text{inr}_2(\Omega)$. The corresponding change in the isoperimetric profile is then $\frac{\text{length}(\partial(\text{inB}_2(\Omega)))}{\text{area}(\text{inB}_2(\Omega))} = \frac{2}{\text{inr}_2(\Omega)}$.

If no disconnected components are generated, a continuous growth in $E_\Omega(t)$ corresponds to increasing g along the medial axis. As derived in equation (7), the slope this induces is $\frac{1}{r(x)}$.

As long as $\frac{1}{r(x)} \leq \frac{2}{\text{inr}_2(\Omega)}$, the IP will not prefer to generate a new disconnected region. While $t < \text{area}(\Omega \circ B_{r_n})$, the largest value of $\frac{1}{r(x)}$ will occur at $\frac{1}{r_n(x)}$. Therefore if $r_n(x) \geq \frac{\text{inr}_2(\Omega)}{2}$, then $E_\Omega(t)$ will always prefer traversing the medial axis.

This proposition excludes cases where $E_\Omega(t)$ could transform discontinuously without increasing number of components and we conjecture that such a change would not be optimal. \square

References

- [AVP99] ARINYO R. J., VIDAL L. P., PASTÓ J. V.: Computing the medial axis transform of polygonal domains by tracing paths. 4
- [BES49] BESICOVITCH A.: A variant of a classical isoperimetric problem. *Quarterly Journal of Mathematics* (01 1949), 84–94. doi: 10.1093/qmath/os-20.1.84. 3
- [Blu67] BLUM H.: A Transformation for Extracting New Descriptors of Shape. In *Models for the Perception of Speech and Visual Form*, Wathen-Dunn W., (Ed.). MIT Press, Cambridge, 1967, pp. 362–380. 4
- [BNNS19] BAR-NATAN A., NAJT L., SCHUTZMAN Z.: The gerrymandering jumble: Map projections permute districts' compactness scores. *arXiv:1905.03173* (2019). 1, 2
- [BS20] BARNES R., SOLOMON J.: Gerrymandering and compactness: implementation flexibility and abuse. *Political Analysis* (2020). 1, 2
- [CDPP16] CHAMBOLLE A., DUVAL V., PEYRÉ G., POON C.: Geometric properties of solutions to the total variation denoising problem. *Inverse Problems* 33, 1 (dec 2016), 015002. 8
- [CFG12] CROFT H., FALCONER K., GUY R.: *Unsolved Problems in Geometry: Unsolved Problems in Intuitive Mathematics*. Problem Books in Mathematics. Springer New York, 2012. URL: <https://books.google.com/books?id=rDTBwAAQBAJ>. 3
- [DLSS19] DEFORD D. R., LAVENANT H., SCHUTZMAN Z., SOLOMON J.: Total variation isoperimetric profiles. *SIAM J. Appl. Algebra Geom.* 3 (2019), 585–613. 2, 8, 9, 10, 12
- [DT18] DUCHIN M., TENNER B. E.: Discrete geometry for electoral geography, 2018. [arXiv:1808.05860](https://arxiv.org/abs/1808.05860). 2
- [Ehr82] EHRENBURG K.: Studien zur hessnung der horizontalen gliederung von erdräuenen. In *Verhandlungen derPhysikalisch-medizinische Gesellschaft zu Würzburg* (1982), p. 41–92. 2
- [FIP20] FIPS: *State FIPS Code Listing*, 2020. URL: https://www.mcc.co.mercer.pa.us/dps/state_fips_code_listing.htm. 9
- [GMT80] GONZALEZ E., MASSARI U., TAMANINI I.: *On the Regularity of Boundaries of Sets Minimizing Perimeter with a Volume Constraint*. Libera Università di Trento. Dipartimento di Matematica, 1980. URL: <https://books.google.com/books?id=jPYcrgEACAAJ>. 2, 3, 6
- [GW96] GORDON C., WEBB D.: You can't hear the shape of a drum. *American Scientist* 84, 1 (1996), 46–55. 10
- [Kar20] KARAVELAS M.: 2D segment Delaunay graphs. In *CGAL User and Reference Manual*, 5.0.2 ed. CGAL Editorial Board, 2020. URL: <https://doc.cgal.org/5.0.2/Manual/packages.html#PkgSegmentDelaunayGraph2>. 5
- [KKKng] KAUFMAN A., KING G., KOMISARCHIK M.: How to measure legislative district compactness if you only know it when you see it. *American Journal of Political Science* (Forthcoming). 2
- [Kno17] KNOPP K.: Einheitliche erzeugung und darstellung der kurven von Peano, Osgood und Koch. *Arch. Math. Phys.* 26 (1917), 103–115. 3
- [Lan77] LANTUÉJOL C.: Sur le modèle de Johnson-Mehl généralisé. In *Internal Report of the Centre de Morph. Math.* (1977). 4
- [LNS17] LEONARDI G., NEUMAYER R., SARACCO G.: The Cheeger constant of a Jordan domain without necks. *Calculus of Variations and Partial Differential Equations* 56 (11 2017), 164. doi:10.1007/s00526-017-1263-0. 3
- [LS19] LEONARDI G. P., SARACCO G.: Minimizers of the prescribed curvature functional in a jordan domain with no necks. *arXiv preprint arXiv:1912.09462* (2019). 1, 2, 3, 9, 12
- [PP91] POLSBY D. D., POPPER R. D.: The third criterion: Compactness as a procedural safeguard against partisan gerrymandering. *Yale Law & Policy Review* 9, 2 (1991), 301–353. 1, 2
- [Ros05] ROS A.: The isoperimetric problem. *Clay Math. Proc.* 2 (01 2005). 2
- [Ser00] SERRA J.: Cours de morphologie mathématique. <http://www.cmm.mines-paristech.fr/~serra/cours/index.htm>, 2000. 3
- [Sul12] SULOVSĀY M.: *Depth, Crossings and Conflicts in Discrete Geometry*. Logos Verlag Berlin, 2012. URL: <https://books.google.com/books?id=0Q9mbXCQRyoc>.
- [SZ98] STREDULINSKY E., ZIEMER W.: Area minimizing sets subject to a volume constraint in a convex set. *Journal of Geometric Analysis* 7 (12 1998). doi:10.1007/BF02921639. 2
- [Vem20] VEMAPS: *Mozambique Map by Vemaps.com*, 2020. URL: <https://vemaps.com/mozambique/mz-01>. 9
- [VF83] VIRGIL., FITZGERALD R.: *The Aeneid*. New York: Random House, 1983. 1
- [voMSRP11] STATE OF MINNESOTA SPECIAL REDISTRICTING PANEL: Order a11-152. http://www.mncourts.gov/Documents/0/Public/Court_Information_Office/2011Redistricting/A110152Order11.4.11.pdf, 11 2011. 2