

Total Variation Isoperimetric Profiles*

Daryl DeFord[†], Hugo Lavenant[‡], Zachary Schutzman[§], and Justin Solomon[¶]

Abstract. Applications such as political redistricting demand quantitative measures of *geometric compactness* to distinguish between simple and contorted shapes. While the isoperimetric quotient, or ratio of area to perimeter squared, is commonly used in practice, it is sensitive to noisy data and irrelevant geographic features like coastline. These issues are addressed in theory by the *isoperimetric profile*, which plots the minimum perimeter needed to inscribe regions of different prescribed areas within the boundary of a shape. Efficient algorithms for computing this profile, however, are not known in practice. Hence, in this paper, we propose a convex Eulerian relaxation of the isoperimetric profile using total variation. We prove theoretical properties of our relaxation, showing that it still satisfies an isoperimetric inequality and yields a convex function of the prescribed area. Furthermore, we provide a discretization of the problem, an optimization technique, and experiments demonstrating the value of our relaxation.

Key words. isoperimetric profile, compactness measures, total variation

AMS subject classifications. 49Q10, 90C90

DOI. 10.1137/18M1215943

1. Introduction and motivation. A classic result in modern geometry, the *isoperimetric inequality* states that the least-perimeter shape enclosing a fixed amount of area is a circle. More formally, if L is the length of a simple closed curve in the plane and A is the area it encloses, then $4\pi A \leq L^2$, with equality if and only if the curve is a circle; the observation that a circle minimizes perimeter subject to a fixed area is a direct byproduct of this expression. Inspired by this inequality, the *isoperimetric quotient* $4\pi A/L^2$ is a commonly used proxy for measuring the compactness of a shape; here, *compactness* refers to geometric regularity, a

*Received by the editors September 21, 2018; accepted for publication (in revised form) August 23, 2019; published electronically November 19, 2019. Authors listed alphabetically.

<https://doi.org/10.1137/18M1215943>

Funding: The fourth author was supported by a Prof. Amar G. Bose Research Grant, the MIT Research Support Committee (“Structured Optimization for Geometric Problems”), Army Research Office grant W911NF-12-R-0011 (“Smooth Modeling of Flows on Graphs”), National Science Foundation grant IIS-1838071 (“BIGDATA:F: Statistical and Computational Optimal Transport for Geometric Data Analysis”), Air Force Office of Scientific Research award FA9550-19-1-0319 (“Structured Assignment: Geometric Optimization Algorithms for Large-Scale Matching”), and an Amazon Research Award. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of these organizations.

[†]Metric Geometry and Gerrymandering Group, Computer Science and Artificial Intelligence Laboratory, Massachusetts Institute of Technology, Cambridge, MA 02139 USA (ddeford@mit.edu).

[‡]Institut de Mathématique d’Orsay, Université Paris-Sud, Orsay Cedex, 91405, and Pacific Institute of Mathematical Sciences, University of British Columbia, Vancouver, BC, Canada, V6T 1Z4 (hugo.lavenant@math.u-psud.fr).

[§]Department of Computer and Information Science, University of Pennsylvania, Philadelphia, PA 19104 USA (ianzach@seas.upenn.edu).

[¶]Corresponding author. Computer Science and Artificial Intelligence Laboratory, Massachusetts Institute of Technology, Cambridge, MA 02139 USA (jsolomon@mit.edu).

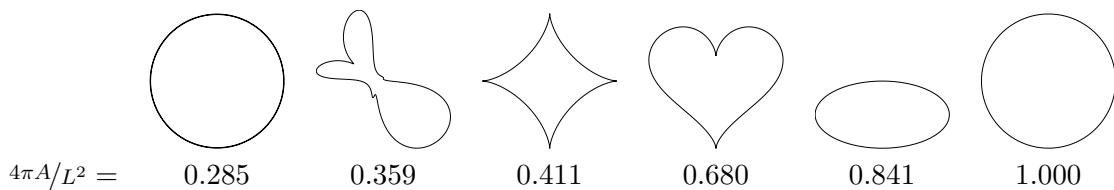


Figure 1. A variety of shapes marked with isoperimetric ratios $4\pi A/L^2$; the least compact shape is a small perturbation of a circle.

piece of terminology that originated in the political geography literature (see section 1.2). This quantity is scale-free, unitless, and bounded between zero and one. It is intended to capture how efficiently a shape “uses” its perimeter to enclose its area and is maximized for a circle.

This quotient, however, is *unstable* in the sense that a small perturbation of the shape’s boundary can greatly increase its perimeter without significantly affecting its area; Figure 1 shows one example of this instability. Moreover, as a measure of compactness, the isoperimetric quotient conflates multiple scales. At a coarse scale, a perturbed circle still appears fairly compact, while at a fine scale, the noisy boundary differentiates it from a proper circle; this distinction cannot be captured by a single number.

To address these issues, in this article we investigate a *multiscale* means of measuring compactness that explicitly assesses aspects of compactness at different scales. Our measure is easily and reliably computed, even on very distorted shapes. Although our work on this problem was inspired by concrete challenges related to the evaluation of compactness of voting districts (see section 1.2), we find that the isoperimetric profile and its convex relaxation have intrinsic mathematical interest that inspire additional challenging problems in geometry and convex analysis.

1.1. Mathematical overview. A modern construction in geometry offers a potential resolution to the instability of the isoperimetric ratio. The *isoperimetric profile* of a shape replaces the isoperimetric quotient with a plot of compactness values at different length or area scales. For each area t between 0 and the area A of a shape ($0 \leq t \leq A$), the isoperimetric profile measures the minimum amount of perimeter *inscribed within* the original shape needed to enclose area t . As shown in Figure 2, by considering $t < A$ we can use this to separate perturbative boundary effects from larger-scale barriers to compactness.

Compared to most of the mathematical literature (see, e.g., [41]), we consider the *full* perimeter of the inscribed shape rather than the relative perimeter; the latter does not account for the perimeter shared with the boundary of the input shape. Considering only relative perimeters—or, equivalently, the isoperimetric problem in domains without boundaries—leads to an elegant mathematical structure with a strong interplay between the geometry of the domain and the shape of the profile. The problem considered in our work, which takes the full perimeter into account, is less studied theoretically (see [45] for some relevant discussion, as well as the forthcoming article [34], which was informed by a preprint of the present article), likely because the shape of the profile has less elegant qualitative properties. This case, however, is relevant for computational problems in need of stable, multiscale compactness measures.

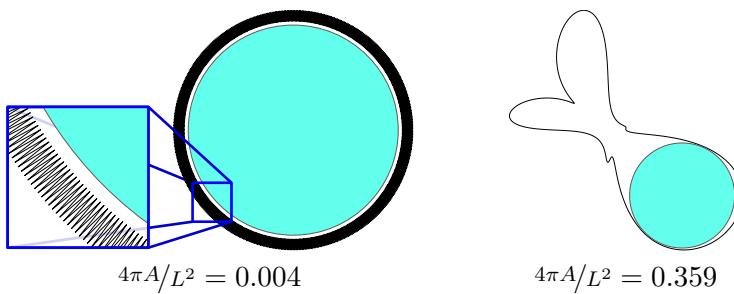


Figure 2. The perturbed circle (left) is not compact as measured by the isoperimetric ratio, but it admits a compact inscribed circle with slightly less area; the noncompact shape on the right does not contain a compact inscribed shape that uses most of the interior area.

Although the isoperimetric profile is a promising theoretical construction, to date no algorithms are known to compute this plot in practice. While special properties in two dimensions suggest that a computational geometry technique *may* be possible to formulate for polygons in the plane, this problem remains open; even less is known about computing the profile in the higher-dimensional case or on Riemannian manifolds. One possible computational approach might use a phase field approximation à la Modica–Mortola [9], but the resulting problem would require the minimization of a nonconvex functional.

Given the potential applications of the isoperimetric profile and the challenges of computing it in practice, this paper proposes a convex relaxation built from total variation. Our basic approach is to write the optimization problem underlying the isoperimetric profile in Eulerian language and then to relax an integer variable to be real-valued. The resulting problem is straightforward to optimize after using standard discretization techniques from mathematical imaging. We show that a key theoretical property of the isoperimetric profile is preserved in our convex relaxation, namely, that the lower envelope of our set of relaxed profiles is provided by a circle—the most compact shape; we also provide theoretical results giving qualitative intuition for the behavior of our profile and its relationship to the original nonconvex problem.

1.2. Compactness and political redistricting. A key application of the isoperimetric quotient—known as the *Polsby–Popper* measure in political science [37]—is in measuring the compactness of legislative voting districts. Here, we use the word *compact*, as used in the political geography literature, to refer to plane regions that appear sufficiently regular, as there is no legally agreed-upon definition of this concept; see [28] for further analysis and discussion.

In representative systems like the U.S. House of Representatives, voters are clustered geographically into districts, with each district electing a single representative to Congress. Manipulating district shapes to engineer a particular outcome for a vote is a practice known as *gerrymandering*, which undermines democratic principles and has been used to deny underrepresented minorities the opportunity to elect a representative of their choosing. While a specific measure of geometric quality is not always mandated by law, scores like *Polsby–Popper* are used as quantitative proxies for the reasonableness of a district or districting plan; contorted, nonconvex district shapes may signal that a district was designed with motives

other than those required by law.

While compactness scores like the isoperimetric quotient are widely used in arguments for or against districting plans, they have significant mathematical drawbacks that undermine their interpretability and reliability [20]. Most importantly, these scores are *unstable*. As shown in Figure 1, the Polsby–Popper score is unstable under boundary perturbations. A visually insignificant adjustment to the boundary of a circle, which has isoperimetric quotient equal to one, can make the isoperimetric quotient arbitrarily close to zero; a minor perturbation changes the shape from being the “most” compact to the “least” compact. This is by no means a degenerate edge-case, but rather extremely common in geographic information systems (GIS) data due to fractal-like structures resulting from the degree of precision with which the geography is measured; maps of different *resolution* of the same district can lead to Polsby–Popper scores that differ by a factor of 2–3 [2]. This particular sensitivity can greatly distort the evaluation of the compactness of a district, since natural choices of boundaries for districts, such as coastlines, geological features, or municipality boundaries, may be noisy. This instability makes it difficult for the isoperimetric quotient to distinguish between districts whose boundaries are contorted due to gerrymandering rather than those that simply contain a coastline or follow a complicated municipal boundary.

The isoperimetric profile, which replaces a single isoperimetric quotient with a whole plot of compactness scores and whose convex relaxation can be computed reliably, can be used to discriminate between compact and noncompact districts without these drawbacks. To demonstrate the interest of this new measurement, in this paper we accompany synthetic experiments with illustrations of the (relaxed) isoperimetric profile of actual U.S. voting districts in the state of North Carolina.

We mention a related notion in the political science literature: A compactness score measuring the proportion of the area of the region filled by its largest inscribed circle is known as the *Ehrenburg test*, presented in a critique of various compactness measures developed in the nineteenth century to describe the geographic and human features of landmasses [21]. This measure captures information about an intermediate scale that depends on the shape of the region being considered but again summarizes compactness with a single number.

2. Background and preliminaries. We begin our discussion with some basic notation and background information setting the stage for our mathematical construction in section 3.

2.1. Isoperimetry. Although the basic isoperimetric problem dates back centuries, the development of isoperimetric inequalities remains an active area of research in mathematics; for a general survey, we refer the reader to [41] for a broad discussion of classical results and open problems. Here, we restrict to basic Euclidean constructions relevant to our computational task.

Take $\Omega \subseteq \mathbb{R}^n$ to be a compact region whose boundary $\partial\Omega \subseteq \Omega$ is an $(n-1)$ -dimensional hypersurface in \mathbb{R}^n , and let B_1 be the unit ball in that space. The general isoperimetric inequality states

$$(2.1) \quad n \cdot \text{vol}(\Omega)^{(n-1)/n} \cdot \text{vol}(B_1)^{1/n} \leq \text{area}(\partial\Omega).$$

This expression encodes the fact that the unit ball B_1 minimizes $(n-1)$ -dimensional boundary surface area over all regions Ω in \mathbb{R}^n with volume 1. As we have already seen, however, it is

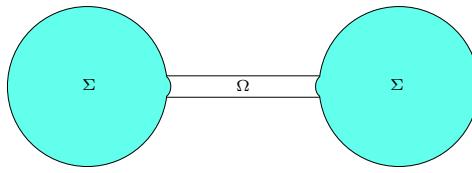


Figure 3. The isoperimetric profile of a nonconvex shape may require a disconnected domain Σ inscribed in Ω ; for example, the most efficient use of area inscribed in the barbell shape above might be disconnected since the rectangular portion in the middle has little area but considerable perimeter. Note this is a schematic rather than computed algorithmically.

straightforward to perturb the boundary of Ω slightly to increase the right-hand side (r.h.s.) of this expression an arbitrary amount with minor impact on the left-hand side (l.h.s.); this is a potential source of instability in practice.

More generally, take $\Omega \subseteq \mathbb{R}^n$ to be the same region, and take $t \in [0, \text{vol}(\Omega)]$. For each t we could ask for the smallest surface area needed to enclose volume t completely within Ω :

$$(2.2) \quad I_\Omega(t) := \min\{\text{area}(\partial\Sigma) : \Sigma \subseteq \Omega \text{ and } \text{vol}(\Sigma) = t\}.$$

Here, we define $\text{area}(\partial\Sigma)$ to be the area of $\partial\Sigma$ as a submanifold of \mathbb{R}^n , that is, including the area of the intersection $\partial\Sigma \cap \partial\Omega$; we put no topological restrictions—in particular, connectedness—on Σ . For our analysis in later sections to be relevant, the minimal regularity assumption needed on Ω is that the TV profile $I_\Omega^{\text{TV}}(t)$ (to be defined in (3.2)) is finite whenever $t < \text{vol}(\Omega)$. This condition can be checked with minimal assumptions on $\partial\Omega$ (including fractality); see Appendix A for a discussion. A direct consequence of the isoperimetric inequality (2.1) is that for any fixed t and prescribed volume V , minimizers of $I_\Omega(t)$ over all possible regions $\Omega \subseteq \mathbb{R}^n$ with $\text{vol}(\Omega) = V$ are those containing a ball with volume t .

In any event, under fairly weak assumptions, a minimizer of problem (2.2) is known to exist. Away from $\partial\Omega$, the boundary $\partial\Sigma$ has constant mean curvature and is nonsingular for $n < 8$ [23]. Because we include the area of $\partial\Sigma \cap \partial\Omega$ in our variational problem, $\partial\Sigma$ and $\partial\Omega$ meet tangentially. The free boundaries of Σ are minimal surfaces (see, e.g., [41, Theorem 1]).

When $n = 2$, the free boundaries of Σ are circular arcs with the same signed radius [4, 5]. This property *may* be helpful in designing a computational geometry-style algorithm for computing $I_\Omega(t)$ when Ω is a polygon in \mathbb{R}^2 , although to our knowledge no such algorithm has been proposed; it remains unknown in which cases computing $I_\Omega(t)$ can be done in polynomial time. As shown in Figure 3, the optimal Σ might be disconnected, even in two dimensions. Additional open problems in the $n = 2$ case are posed in [19]. To make matters more challenging, this helpful structure does not appear to extend for $n \geq 3$, underscoring the potential value of our construction in section 3 even if I_Ω is efficiently computable in some special cases.

Open Problem 2.1. Identify a polynomial-time algorithm or NP-hardness result for computing isoperimetric profiles. The simplest open problem is computing the isoperimetric profile of a polygon in the plane \mathbb{R}^2 ; if evaluating the profile of such a polygon is polynomial-time solvable, higher-dimensional analogues are of interest as well.

2.2. Total variation. The definition of the isoperimetric profile (2.2) is Lagrangian: The variable Σ explicitly parameterizes the boundary of an unknown shape. Our approach in

section 3 will be to switch an Eulerian formulation, replacing the boundary shape optimization variable with the indicator function of the unknown Σ . To do so, we need an Eulerian way to compute the area of $\partial\Sigma$, which we can achieve using the *total variation* of its indicator.

For a function $f \in L^1(\mathbb{R}^n)$, the total variation (TV) of f is defined as

$$(2.3) \quad \text{TV}[f] := \sup \left\{ \int_{\mathbb{R}^n} [f(x) \nabla \cdot \phi(x)] dx : \phi \in C_c^1(\mathbb{R}^n \rightarrow \mathbb{R}^n) \text{ and } \|\phi\|_\infty \leq 1 \right\}.$$

Here, $C_c^1(\mathbb{R}^n \rightarrow \mathbb{R}^n)$ denotes the space of compactly supported continuously differentiable functions from \mathbb{R}^n into \mathbb{R}^n . For differentiable functions $f \in C^1(\mathbb{R}^n)$ with compact support, the divergence theorem implies an alternative formula for total variation introduced to mathematical image processing in [43]:

$$(2.4) \quad \text{TV}[f] = \int_{\mathbb{R}^n} \|\nabla f\|_2 dx.$$

An alternative way of computing total variation, which reveals its link with the isoperimetric problem, is the coarea formula [22]. For a nonnegative f , this formula states

$$(2.5) \quad \text{TV}[f] = \int_0^{+\infty} \text{area}(\partial\{f \geq s\}) ds.$$

For a region $\Sigma \subseteq \mathbb{R}^n$, denote its *indicator function* $\mathbb{1}_\Sigma$ via

$$(2.6) \quad \mathbb{1}_\Sigma(x) := \begin{cases} 1 & \text{if } x \in \Sigma, \\ 0 & \text{otherwise.} \end{cases}$$

Then a consequence of the coarea formula (2.5) is that

$$(2.7) \quad \text{area}(\partial\Sigma) = \text{TV}[\mathbb{1}_\Sigma].$$

A detailed account of geometric information measured by total variation is provided in [14].

3. Total variation isoperimetric profile. Now that we have established basic notation and the functionals we will use in our construction, we are ready to present our convex relaxation of the isoperimetric profile for a shape embedded in \mathbb{R}^n .

3.1. Definition. Inspired by the total variation formula (2.7), we can formally rewrite the optimization problem (2.2) for the isoperimetric profile in an Eulerian fashion:

$$(3.1) \quad I_\Omega(t) = \begin{cases} \inf_{f \in L^1(\mathbb{R}^n)} \text{TV}[f] \\ \text{subject to } \int_{\mathbb{R}^n} f(x) dx = t, \\ 0 \leq f \leq \mathbb{1}_\Omega, \\ f(x) \in \{0, 1\} \forall x \in \mathbb{R}^n. \end{cases}$$

Here, f is the indicator of the unknown shape Σ . The three constraints ensure that (1) the area enclosed in Σ is equal to t , (2) Σ is inscribed within Ω , and (3) f is properly an indicator function. For convenience, we have written this problem using integrals over \mathbb{R}^n , but in reality the second constraint ensures that f is zero outside Ω .

Since total variation is convex, the form (3.1) hides all the nonconvexity of the problem in the third constraint. Hence, we propose using an alternative measure of compactness that drops the nonconvex constraint, which we call the *total variation (TV) profile* of Ω :

$$(3.2) \quad I_{\Omega}^{\text{TV}}(t) := \begin{cases} \min_{f \in L^1(\mathbb{R}^n)} \text{TV}[f] \\ \text{subject to} \quad \int_{\mathbb{R}^n} f(x) dx = t, \\ \quad 0 \leq f \leq \mathbb{1}_{\Omega}. \end{cases}$$

The existence of a minimizer is an immediate consequence of Rellich's compactness theorem [1, Theorem 3.23] (also see [13, Theorem 2]).

Since (3.2) is obtained by dropping a constraint from (3.1), we immediately have the following.

Proposition 3.1. *For all (Ω, t) , we have $I_{\Omega}^{\text{TV}}(t) \leq I_{\Omega}(t)$.*

Unfortunately, this inequality is not tight, which we demonstrate in the following example.

Example 3.2 (circle). Suppose $\Omega \subset \mathbb{R}^2$ is a circle of radius R , and take $t = \pi r^2$ for $r \in (0, R)$. In this case, by the isoperimetric inequality we know $I_{\Omega}(t) = 2\pi r$. But suppose we take $f(x) \equiv r^2/R^2$. Notice $f(x)$ satisfies the constraints in (3.2), but by the coarea formula

$$I_{\Omega}^{\text{TV}}(t) \leq \text{TV}[f] = 2\pi R \cdot \frac{r^2}{R^2} = 2\pi r \cdot \frac{r}{R} < I_{\Omega}(t).$$

Hence, our relaxation is not tight.

As we will prove in Proposition 3.5, the TV profile still admits an elegant characterization explaining why the relaxation is not tight: it is the lower convex envelope of the isoperimetric profile.

The following monotonicity property can be derived directly from the definition of the total variation profile.

Proposition 3.3. *Suppose $\Omega_1 \subseteq \Omega_2 \subset \mathbb{R}^d$ are two compact domains. Then for all $t \leq \text{vol}(\Omega_1)$ it holds that $I_{\Omega_2}^{\text{TV}}(t) \leq I_{\Omega_1}^{\text{TV}}(t)$.*

In particular, if a domain Ω can be squeezed between two other domains, $\Omega_1 \subseteq \Omega \subseteq \Omega_2$, then the TV profile of Ω is squeezed between those of Ω_1 and Ω_2 , at least for $t \leq \text{vol}(\Omega_1)$. For instance, the TV profile of the perturbed circle (Figure 2, left) must be very close, at least for t away from $\text{vol}(\Omega)$, to a straight line, the TV profile of a circle. More generally, if we know the domain Ω with some uncertainty or only at a given resolution, which is roughly equivalent to saying that we can squeeze it between two other domains with a controlled difference in area, then we have guarantees for the value of the TV profile as soon as t is not too close to $\text{vol}(\Omega)$.

3.2. Duality and convexity. We next derive a dual for the TV profile problem (3.2) to provide additional insight into the structure of the problem and in particular to show that $I_{\Omega}^{\text{TV}}(t)$ is convex in t . We begin by writing a minimax formulation of the problem by

substituting (2.3):

$$(3.3) \quad I_{\Omega}^{\text{TV}}(t) = \min_{f \in L^1(\mathbb{R}^n)} \left\{ \begin{array}{l} \sup_{\phi, \lambda, \psi, \xi} \int_{\mathbb{R}^n} [f(x) \nabla \cdot \phi(x)] dx \\ \quad + \lambda \left(t - \int_{\mathbb{R}^n} f(x) dx \right) \\ \quad - \int_{\mathbb{R}^n} [\psi(x) f(x) + \xi(x) (\mathbb{1}_{\Omega}(x) - f(x))] dx \\ \text{subject to} \quad \|\phi\|_{\infty} \leq 1, \psi \geq 0, \xi \geq 0. \end{array} \right.$$

Here, the dual variables are $\phi \in C_c^1(\mathbb{R}^n \rightarrow \mathbb{R}^n)$, $\lambda \in \mathbb{R}$, and $\psi, \xi \in C_c^1(\mathbb{R}^n \rightarrow \mathbb{R})$.

A standard argument using the Fenchel–Rockafellar duality theorem [11, Theorem 1.12] justifies swapping the inner and outer problems; for completeness, a formal argument is in Appendix B. Removing extraneous terms, we are left with the following inner variational problem for f , with the dual variables fixed:

$$\inf_{f \in L^1(\mathbb{R}^n)} \int_{\mathbb{R}^n} f(x) [\nabla \cdot \phi(x) - \lambda - \psi(x) + \xi(x)] dx.$$

This problem is unbounded unless $\nabla \cdot \phi(x) + \xi(x) - \psi(x) = \lambda$ a.e. $x \in \mathbb{R}^n$. Hence, after swapping the inner and outer problems in (3.3) and substituting this relationship we find

$$I_{\Omega}^{\text{TV}}(t) = \left\{ \begin{array}{l} \sup_{\phi, \lambda, \psi, \xi} \lambda t - \int_{\Omega} \xi(x) dx \\ \text{subject to} \quad \|\phi\|_{\infty} \leq 1, \psi \geq 0, \xi \geq 0, \\ \quad \nabla \cdot \phi(x) + \xi(x) - \psi(x) = \lambda \text{ a.e. } x \in \mathbb{R}^n. \end{array} \right.$$

Define $\eta(x) := \xi(x) - \psi(x)$. From the objective of the problem above, we can see $\xi(x) = \max(\eta(x), 0)$ and $\psi(x) = \max(-\eta(x), 0)$. Substituting and simplifying leads to the dual

$$(3.4) \quad I_{\Omega}^{\text{TV}}(t) = \left\{ \begin{array}{l} \sup_{\phi \in C_c^1(\mathbb{R}^n \rightarrow \mathbb{R}^n), \lambda \in \mathbb{R}} \lambda t - \int_{\Omega} \max(\lambda - \nabla \cdot \phi(x), 0) dx \\ \text{subject to} \quad \|\phi\|_{\infty} \leq 1 \end{array} \right.$$

While this dual formula provides an interesting—if abstract—reinterpretation of $I_{\Omega}^{\text{TV}}(t)$ in its own right, it also allows us to derive another property of the function itself.

Proposition 3.4. $I_{\Omega}^{\text{TV}}(t)$ is a convex function of t .

Proof. Define

$$h(\lambda) := \inf_{\|\phi\|_{\infty} \leq 1} \int_{\Omega} \max(\lambda - \nabla \cdot \phi(x), 0) dx.$$

From (3.4), we see that $I_{\Omega}^{\text{TV}}(t) = h^*(t)$, the convex conjugate (Legendre–Fenchel transform) of h , which is necessarily convex [6]. \blacksquare

We can say even more: The TV profile I_{Ω}^{TV} is the lower convex envelope of the isoperimetric profile I_{Ω} , i.e., the largest convex function upper-bounded by the isoperimetric profile.

Proposition 3.5. The function I_{Ω}^{TV} is the lower convex envelope of I_{Ω} .

Proof. It is known that the lower convex envelope of a function is equal to the Legendre transform of its Legendre transform [40, section 12]. Since Proposition 3.4 shows that I_{Ω}^{TV} is convex—and thus equals its convex envelope—it suffices to prove that the Legendre transform of I_{Ω} coincides with that of I_{Ω}^{TV} .

By definition,

$$(3.5) \quad (I_\Omega)^*(\lambda) = \sup_t [\lambda t - I_\Omega(t)] = \begin{cases} \sup_f \lambda \int_\Omega f(x) dx - \text{TV}(f) \\ \text{subject to } 0 \leq f \leq \mathbb{1}_\Omega, \\ f(x) \in \{0, 1\} \forall x \in \mathbb{R}^n. \end{cases}$$

On the other hand, the Legendre transform of I_Ω^{TV} is given by

$$(3.6) \quad (I_\Omega^{\text{TV}})^*(\lambda) = \sup_t \lambda t - I_\Omega(t) = \begin{cases} \sup_f \lambda \int_\Omega f(x) dx - \text{TV}(f) \\ \text{subject to } 0 \leq f \leq \mathbb{1}_\Omega. \end{cases}$$

Since there are more admissible competitors in (3.6) than (3.5), $(I_\Omega^{\text{TV}})^*(\lambda) \geq (I_\Omega)^*(\lambda)$. On the other hand, using the coarea formula (2.5), for any admissible competitor f in (3.6),

$$\lambda \int_\Omega f(x) dx - \text{TV}(f) = \int_0^1 \underbrace{(\lambda \text{vol}(\{f \geq s\}) - \text{area}(\partial\{f \geq s\}))}_{\leq (I_\Omega)^*(\lambda) \text{ by definition}} ds \leq (I_\Omega)^*(\lambda).$$

Taking the supremum in f leads to $(I_\Omega^{\text{TV}})^*(\lambda) \leq (I_\Omega)^*(\lambda)$, which concludes the proof. \blacksquare

The proof of the equality of (3.5) and (3.6) is just a slight adaptation of [13, Proposition 2.1], in a simpler case, since we want only the equality of the values of the minimization problem rather than a correspondence between the minimizers.

3.3. Behavior for small $t > 0$. Continuing in our effort to describe the shape of $I_\Omega^{\text{TV}}(t)$ as precisely as possible, we can accompany the convexity result in Proposition 3.4 with a description of its behavior when t is close to zero. In particular, we will show that I_Ω^{TV} is *linear* with positive slope for small t , a property we can verify in our experiments. This is not surprising, given Proposition 3.5, since the isoperimetric profile of Ω coincides with that of \mathbb{R}^n for small t , the latter being concave. Nevertheless, we can say more by describing exactly the slope of I_Ω^{TV} for small volumes.

We use the following auxiliary problem, known as the Cheeger problem [36]. The Cheeger constant of Ω , denoted by $h_1(\Omega)$, is defined as

$$h_1(\Omega) := \inf_{\tilde{\Sigma} \subseteq \Omega} \frac{\text{area}(\partial\tilde{\Sigma})}{\text{vol}(\tilde{\Sigma})},$$

and a subset $\Sigma \subseteq \Omega$ such that $h_1(\Omega) = \frac{\text{area}(\partial\Sigma)}{\text{vol}(\Sigma)}$ is known as a Cheeger set of Ω . Such a set exists as soon as Ω has a Lipschitz boundary [36, Proposition 3.1] and is unique if Ω is convex [36, Proposition 5.2]. Σ can be interpreted as the largest “smooth” subset of Ω . An explicit description of Cheeger sets can be found for convex domains [29], as well as domains without necks [32]; we emphasize, however, that domains appearing in redistricting applications are not likely to satisfy the assumptions of these articles.

Furthermore, provided we know a Cheeger set of Ω , we can describe the behavior of $I_\Omega^{\text{TV}}(t)$ for small t .

Proposition 3.6. *Let Ω be compact, let $h_1(\Omega)$ be the Cheeger constant of Ω , and let Σ be a Cheeger set of Ω . Then for any $t \leq \text{vol}(\Sigma)$, we have $I_\Omega^{\text{TV}}(t) = h_1(\Omega)t$, and a solution f in (3.2) is given by $f := \frac{t}{\text{vol}(\Sigma)} \cdot \mathbb{1}_\Sigma$.*

Proof. It is clear that $\hat{f} := \frac{t}{\text{vol}(\Sigma)} \cdot \mathbb{1}_\Sigma$ satisfies the constraints of problem (3.2) defining $I_\Omega^{\text{TV}}(t)$ as soon as $t \leq \text{vol}(\Sigma)$, which ensures $0 \leq \hat{f} \leq \mathbb{1}_\Sigma \leq \mathbb{1}_\Omega$. Hence, $I_\Omega^{\text{TV}}(t) \leq h_1(\Omega)t$. On the other hand, using the coarea formula (2.5), if f is any competitor for problem (3.2), then

$$\begin{aligned} \text{TV}(f) &= \int_0^{+\infty} \text{area}(\partial\{f \geq s\}) ds \\ &= \int_0^{+\infty} \text{vol}(\{f \geq s\}) \cdot \underbrace{\frac{\text{area}(\partial\{f \geq s\})}{\text{vol}(\{f \geq s\})}}_{\geq h_1(\Omega) \text{ by definition}} ds \\ &\geq h_1(\Omega) \int_0^{+\infty} \text{vol}(\{f \geq s\}) ds = h_1(\Omega) \int_{\mathbb{R}^d} f(x) dx = h_1(\Omega)t. \end{aligned}$$

Hence, for $t \leq \text{vol}(\Omega)$, we have $I_\Omega^{\text{TV}}(t) = h_1(\Omega)t$. ■

If Ω is not convex, then there may exist more than one Cheeger set in Ω [33, section 4]. In particular, in light of the proof of Proposition 3.6, there may be more than one minimizer of (3.2). Since we use an interior point solver in our experiments, the presence of multiple solutions for nonconvex boundaries explains why our solutions appear fuzzy in this case.

Considering Proposition 3.6, a byproduct of our numerical method detailed below is the ability to compute Cheeger constants and Cheeger sets, at least when they are unique. This problem was tackled numerically in [12], where the authors propose a method to compute the largest Cheeger set of a given set Ω . Compared to our numerical results, theirs are sharper; they also can characterize precisely which set is selected. On the other hand, we are largely interested in the optimal objective *value*, i.e., the Cheeger constant, which does not depend on the fuzziness of the minimizer f ; moreover, the link to the Cheeger problem—and hence the numerical method of [12]—is lost as soon as t becomes large. We leave as future work a detailed exploration of the relationship between these two techniques.

3.4. Structure of the minimizers. As we have seen in the previous section, at least for small t one can choose a solution f of the problem (3.2) that is proportional to an indicator function, i.e., that takes only two values. Due to the potential nonuniqueness of Cheeger sets, however, it is possible to construct optimal functions f that take on infinitely many values, even for small t . Up to the selection of appropriate optimal minimizers for the TV profile, however, we can prove the following.

Proposition 3.7. *There exists a family $(f_t)_{t \in [0,1]}$ such that the following hold:*

- For any $t \in [0, 1]$, the function $f_t \in L^1(\mathbb{R}^n)$ satisfies $0 \leq f_t \leq \mathbb{1}_\Omega$, $\int_{\mathbb{R}^n} f_t(x) dx = t$ and $\text{TV}(f_t) = I_\Omega^{\text{TV}}(t)$.
- For any $t \in [0, 1]$, there exist $v_t \in (0, 1)$ such that f_t takes its values in $\{0, v_t, 1\}$.
- For a.e. $x \in \Omega$, the function $t \rightarrow f_t(x)$ is increasing.

Proof. Let $S \subseteq [0, \text{vol}(\Omega)]$ be the union of $\{0, \text{vol}(\Omega)\}$ and the set of t such that I_Ω^{TV} is not an affine function in a neighborhood of t ; the latter is closed in $[0, \text{vol}(\Omega)]$. Since I_Ω is lower semicontinuous and I_Ω^{TV} is its lower convex envelope, for any $t \in [0, \text{vol}(\Omega)]$, if $I_\Omega^{\text{TV}}(t) < I_\Omega(t)$, then I_Ω^{TV} is affine in a neighborhood of t . In particular, $I_\Omega^{\text{TV}}(t) = I_\Omega(t)$ for any point $t \in S$.

Let $t \leq s$ with $t, s \in S$. By the previous remark, there exist subsets Σ_t and Σ_s of Ω , with respective volumes t and s , such that $\text{area}(\partial\Sigma_t) = I_\Omega(t) = I_\Omega^{\text{TV}}(t) = \text{TV}(\mathbb{1}_{\Sigma_t})$; a similar identity holds for Σ_s . Let us consider $\Sigma_+ := \Sigma_t \cup \Sigma_s$ and $\Sigma_- := \Sigma_t \cap \Sigma_s$; we denote $r := \text{vol}(\Sigma_-)$. By the inclusion–exclusion principle, $\text{vol}(\Sigma_+) = t + s - r$. On the other hand, using [1, Proposition 3.38], the perimeter of Σ_+ satisfies

$$\text{area}(\partial\Sigma_+) \leq \text{area}(\partial\Sigma_t) + \text{area}(\partial\Sigma_s) - \text{area}(\partial\Sigma_-).$$

Using $\mathbb{1}_{\Sigma_+}$ as a competitor for the problem (3.2) defining $I_\Omega^{\text{TV}}(t + s - r)$, and given that $\text{area}(\partial\Sigma_-) \geq I_\Omega^{\text{TV}}(r)$, we deduce that

$$(3.7) \quad I_\Omega^{\text{TV}}(t + s - r) \leq I_\Omega^{\text{TV}}(t) + I_\Omega^{\text{TV}}(s) - I_\Omega^{\text{TV}}(r).$$

Convex nonnegative functions are superadditive, but the equation above indicates that I_Ω^{TV} exhibits subadditive behavior. More precisely, using Lemma C.1 proved in the appendix, we see that I_Ω^{TV} is affine on $[r, t + s - r]$ as soon as $r < t$. Since $t, s \in S$, the only way to avoid a contradiction is if $r = t$, which implies $\Sigma_t \subseteq \Sigma_s$.

As a consequence, using $t = s$, we see that for any $t \in S$ there exists a unique $\Sigma_t \subseteq \Omega$ such that $I_\Omega^{\text{TV}}(t) = \text{TV}(\mathbb{1}_{\Sigma_t})$. Using $t \leq s$ we deduce that the map $t \in S \rightarrow \Sigma_t$ is increasing with respect to (w.r.t.) inclusion. In particular, for a.e. $x \in \mathbb{R}^n$, the map $t \in S \rightarrow \mathbb{1}_{\Sigma_t}(x)$ is increasing.¹ For $t \in S$, we set $f_t := \mathbb{1}_{\Sigma_t}$; such a choice satisfies all the requirements of Proposition 3.7.

Since the set $[0, \text{vol}(\Omega)] \setminus S$ is open, it can be decomposed as a countable union of open intervals. Let (t, s) be one of such interval, i.e., $t, s \in S$ and $S \cap (t, s) = \emptyset$. The definition of S , helped by a connectivity argument and the continuity of I_Ω^{TV} , shows that I_Ω^{TV} is affine on the segment $[t, s]$. On the other hand, for any $r \in [t, s]$, the function

$$(3.8) \quad f_r := \frac{s-r}{s-t} \mathbb{1}_{\Sigma_t} + \frac{r-t}{s-t} \mathbb{1}_{\Sigma_s},$$

i.e., the convex combination of $\mathbb{1}_{\Sigma_t}$ and $\mathbb{1}_{\Sigma_s}$ whose total mass is r , satisfies $\int_\Omega f_r(x) dx = r$, $0 \leq f_r \leq \mathbb{1}_\Omega$ and, by convexity of the TV norm,

$$\text{TV}(f_r) \leq \frac{s-r}{s-t} I_\Omega^{\text{TV}}(t) + \frac{r-t}{s-t} I_\Omega^{\text{TV}}(s).$$

The r.h.s. of this equation is precisely the affine function joining $(t, I_\Omega^{\text{TV}}(t))$ to $(s, I_\Omega^{\text{TV}}(s))$, which implies that the inequality is an equality. Hence, the function f_r satisfies $0 \leq f_r \leq \mathbb{1}_\Omega$, $\int_{\mathbb{R}^n} f_r(x) dx = t$, and $\text{TV}(f_r) = I_\Omega^{\text{TV}}(r)$.

As $\Sigma_t \subset \Sigma_s$, the map $r \in [t, s] \rightarrow f_r(x)$ is increasing for a.e. $x \in \mathbb{R}^n$. Moreover, the function f_r takes the value 0 on $\Omega \setminus \Sigma_s$, the value 1 on Σ_t , and only a third value, namely, $(r-t)/(s-t)$ on the set $\Sigma_s \setminus \Sigma_t$. In conclusion, on every open interval within $[0, \text{vol}(\Omega)] \setminus S$ we define f_r by (3.8), and such a choice satisfies all the requirement of Proposition 3.7. ■

¹For this to hold we have to choose a precise representative of Σ . For instance, we choose the set of points in \mathbb{R}^n such that the Lebesgue density of Σ is equal to 1.

3.5. Isoperimetric inequality. So far, we have provided some propositions describing the shape of I_Ω^{TV} ; our next task is to verify that it has properties in common with the isoperimetric profile that make it useful for evaluating compactness. Although our relaxation does not always recover the solution to the original problem, the following key property is preserved.

Proposition 3.8 (isoperimetry). *Suppose $B \subset \mathbb{R}^n$ is a ball whose volume matches $\text{vol}(\Omega)$. Then, for all $t \in [0, \text{vol}(\Omega)]$, we have $I_B^{\text{TV}}(t) \leq I_\Omega^{\text{TV}}(t)$, and if the equality holds for some $t > 0$, then Ω is a ball.*

Proof. As recalled in the introduction, we already know that $I_B(t) \leq I_\Omega(t)$ for any t . Hence, any convex function bounded by I_B is also bounded by I_Ω . Taking the supremum and recalling Proposition 3.5, we see that $I_B^{\text{TV}}(t) \leq I_\Omega^{\text{TV}}(t)$.

Now assume that $I_B^{\text{TV}}(t) = I_\Omega^{\text{TV}}(t)$ for some $t \in (0, \text{vol}(\Omega)]$. Taking into account the linearity of I_B^{TV} (Example 3.2) and the convexity of I_Ω^{TV} , we know that the slope of I_B^{TV} and I_Ω^{TV} coincide at $t = 0$. Considering Proposition 3.6, this implies that the Cheeger constants of B and Ω agree, which can only happen if Ω is a ball [36, Proposition 6.11]. ■

4. Discretization and optimization. Having established theoretical properties of the TV isoperimetric profile, our next step is to provide a discretization and algorithm for its approximation in practice.

4.1. Discretization. In our experiments, we assume that the shape Ω is expressed as an indicator function on a uniform grid; for example, when $n = 2$ we take as input an image I with $I_{i,j} = 1$ inside Ω and $I_{i,j} = 0$ outside. We use the fourfold discretization ∇_p of the gradient operator ∇ proposed in [14] to promote rotational invariance for our model. Namely, denote by P the total number of grid points and by Δx the grid size. Then the linear operator $\nabla_p : \mathbb{R}^P \rightarrow \mathbb{R}^{P \times 4}$ is defined at any point $p = (i, j)$ as

$$(4.1) \quad (\nabla_p f)_{i,j} = \frac{1}{\Delta x} \begin{pmatrix} f_{i+1,j} - f_{i,j} \\ f_{i,j+1} - f_{i,j} \\ f_{i+1,j+1} - f_{i,j+1} \\ f_{i+1,j+1} - f_{i+1,j} \end{pmatrix}.$$

Using this operator, the total variation of a function defined on the uniform grid is approximated by

$$\sum_p \|\nabla_p f\|_2 = \sum_{i,j} \|(\nabla_p f)_{i,j}\|_2,$$

where $\|\cdot\|_2$ is the Euclidean norm in \mathbb{R}^4 . After padding the grid with zeros on all four sides, we do not need to account for boundary conditions.

Our discretization of (3.2) becomes

$$(4.2) \quad I_\Omega^{\text{discretized}}(t) := \begin{cases} \inf_{f \in \mathbb{R}^P} \sum_p \|\nabla_p f\|_2 \\ \text{subject to } \mathbf{1}^\top f = t, \\ 0 \leq f \leq I. \end{cases}$$

While careful proof of convergence in the limit of grid refinement is outside the scope of our discussion, we note that several analogous results exist for image processing models with similar structure [3, 30].

4.2. Optimization. Our problem (4.2) is a second-order cone program (SOCP) [8], for which there exist extremely efficient industrial solvers. We find that MOSEK [35]—easily called using the CVX library [24, 25]—is effective for up to medium-scale instances and competitive with hand-designed algorithms. For scalability and simplicity, however, we can also derive a first-order algorithm based on the alternating direction method of multipliers (ADMM) [7], detailed below.

Suppose $z \in [0, 1]^{P'}$ is the restriction of the unknown variable f in (4.2) to those pixels p in the image I that are nonzero; the remaining entries of f must equal zero. We can write an equivalent formulation of problem (4.2) as

$$(4.3) \quad I_{\Omega}^{\text{discretized}}(t) := \begin{cases} \inf_{z \in \mathbb{R}^{P'}} \sum_p \|\mathcal{P}_p G z\|_2 \\ \text{subject to } \mathbb{1}^T z = t, \\ \quad 0 \leq z \leq 1. \end{cases}$$

Here, $G \in \mathbb{R}^{4P \times P'}$ is the restriction of the gradient operator ∇ in (4.1) to the nonzero pixels, and $\mathcal{P}_p \in \{0, 1\}^{4 \times 4P}$ extracts the four elements of the gradient relevant to pixel p .

To derive tractable ADMM iterations, we rewrite (4.3) in a somewhat counterintuitive form that leads to a tractable splitting:

$$(4.4) \quad I_{\Omega}^{\text{discretized}}(t) := \begin{cases} \inf_{x, z, z'} \sum_p \|\mathcal{P}_p x\|_2 + \chi(0 \leq z' \leq 1) \\ \text{subject to } x - G z = 0, \quad : y \\ \quad \mathbb{1}^T z = t, \quad : \lambda \\ \quad z - z' = 0. \quad : q \end{cases}$$

This expression copies z into a second variable z' and isolates the gradient Gz as a third variable x ; χ denotes a (convex) indicator function that equals ∞ any time the constraint is violated, and 0 otherwise. The augmented Lagrangian of this optimization problem is

$$\begin{aligned} \Lambda(x, z, z'; y, \lambda, q) := & \sum_p \|\mathcal{P}_p x\|_2 + \chi(0 \leq z' \leq 1) + \frac{\rho}{2} \|x - G z\|_2^2 + y^T(x - G z) \\ & + \frac{\tau}{2} (\mathbb{1}^T z - t)^2 + \lambda(\mathbb{1}^T z - t) + \frac{\beta}{2} \|z - z'\|_2^2 + q^T(z - z'). \end{aligned}$$

Our two-block ADMM scheme cycles between three steps:

$$\begin{aligned} (\text{Primal 1}) \quad & z \leftarrow \min_z \Lambda(x, z, z'; y, \lambda, q) && (\text{linear system}), \\ (\text{Primal 2}) \quad & (x, z') \leftarrow \min_{(x, z')} \Lambda(x, z, z'; y, \lambda, q) && (\text{closed form}), \\ (\text{Dual}) \quad & y \leftarrow y + \rho(x - G z) && (\text{closed form}), \\ & \lambda \leftarrow \lambda + \tau(\mathbb{1}^T z - t), \\ & q \leftarrow q + \beta(z - z'). \end{aligned}$$

Iteratively cycling these three steps is guaranteed to converge to the global optimum, as justified by the discussion in [7, section 3.2.1]. Our splitting is carefully designed to make each iteration above computationally tractable; z is obtained by solving a linear system of

Algorithm 4.1. ADMM TO SOLVE PROBLEM (4.4).

```

function TVPROFILEADMM
   $M \leftarrow \rho G^\top G + \tau \mathbb{1} \mathbb{1}^\top + \beta I_{P' \times P'}$   $\triangleright$  Can be factored once before iteration

  for  $i = 1, 2, \dots$  do
    // Primal 1:  $z$ 
     $r \leftarrow G^\top (\rho x + y) + (t\tau - \lambda) \mathbb{1} + (\beta z' - q)$   $\triangleright$  Right-hand side of linear system
     $z \leftarrow M^{-1}r$   $\triangleright$  Can be accelerated using Cholesky factorization

    // Primal 2:  $(x, z')$   $\triangleright$  Decouples over  $x$  and  $z'$ 
    // Primal 2a:  $x$ 
     $H \leftarrow \text{RESHAPE}(Gz - y/\rho, 4 \times P)$   $\triangleright$  Isolate individual gradient vectors
    for  $p = 1, 2, \dots, P$  do  $\triangleright$  Per-column operation
       $c_p \leftarrow \text{MAX}(1 - 1/\rho \|H_p\|_2, 0)$   $\triangleright H_p$  is column  $p$  of  $H$ 
       $X_p \leftarrow c_p H_p$ 
     $x \leftarrow \text{RESHAPE}(X, 4P \times 1)$   $\triangleright$  Unroll back to a vector

    // Primal 2b:  $z'$ 
     $z' \leftarrow \text{CLAMP}(z + q/\beta, [0, 1])$   $\triangleright$  Dual ascent

    // Dual:  $(y, \lambda, q)$ 
     $y \leftarrow y + \rho(x - Gz)$ 
     $\lambda \leftarrow \lambda + \tau(\mathbb{1}^\top z - t)$ 
     $q \leftarrow q + \beta(z - z')$ 

    // Check for convergence and update  $(\rho, \tau, \beta)$  here  $\triangleright$  See [7, sect. 3.3, sect. 3.4] for details
  return  $(x, z, z')$ 
end function

```

equations, while the (x, z') step decouples over these two variables and is solvable in closed form (soft thresholding and clamping). The linear system for z has the same matrix in each iteration, which can be prefactored. Algorithm 4.1 fills in the details of the steps.

We found this iterative technique to be the most efficient among the many possible methods for convex problems in the form (4.2); for instance, one alternative might be to use proximal splitting as described in [13, section 3.2.4] and [15, section 5]. It is worth noting, however, that we are in the worst-case situation for these algorithms: the terms in the most obvious proximal splitting are neither smooth nor strictly convex. Indeed, at least in the continuous case, the form of the optimal f described in Proposition 3.7 indicates that typical solutions of our problem are at points of nondifferentiability.

Figure 4 tests the efficiency of the ADMM algorithm in section 4.2 on small-scale and large-scale examples. Our experiments involve computation of evenly spaced samples from the TV profile of a nonconvex shape; the size of the image and number P' of unknowns in the convex optimization problem are shown below the plots. Our experiments were carried out in MATLAB 2018b on an Ubuntu 16.04 machine with 64 GB memory and an Intel Xeon Gold 6136 CPU (3.00 GHz).

For each example, we plot the error of the ADMM variable z relative to ground truth computed using MOSEK [35] with relatively high precision ($\epsilon^{3/4}$, where $\epsilon \approx 2.22 \times 10^{-16}$ is machine precision); the relative error of the duplicated variable z' converges similarly. While it is difficult to match MOSEK's and ADMM's convergence criteria exactly, as a point of reference the ground truth computations in MOSEK took an average of 0.726 seconds per small-scale example and 265.5 seconds per large-scale example to converge.

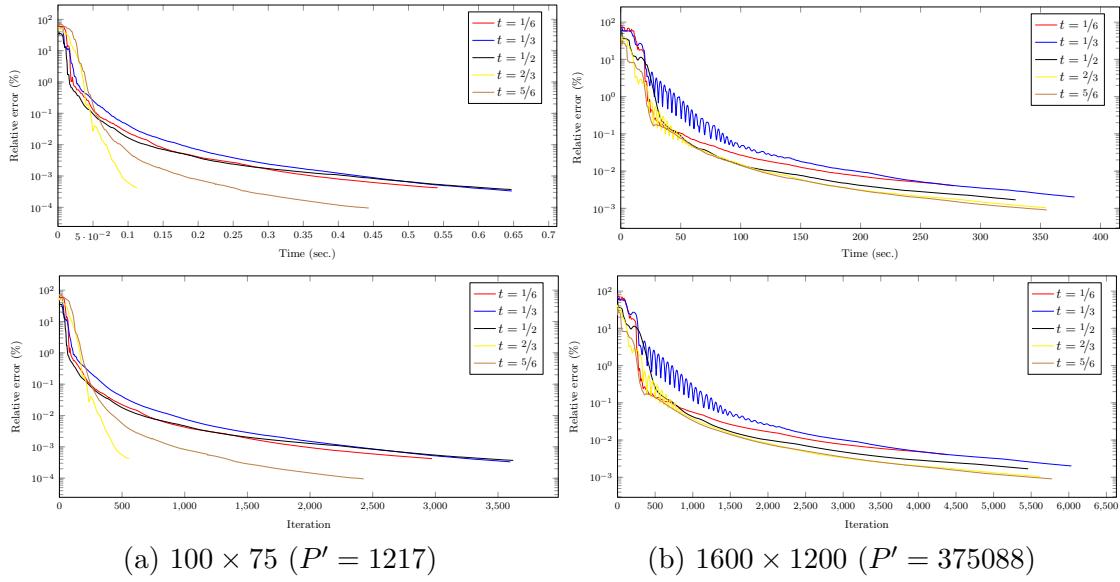


Figure 4. Error relative to ground truth for the ADMM algorithm presented in section 4.2, as a function of time (top) and iteration (bottom). The image used for testing is district 12 in Figure 8. The different tests correspond to evenly spaced t values, $t \in \{1/6, \dots, 5/6\}$ for total area $A = 1$.

Our ADMM implementation is far from optimized, but it does include a few straightforward improvements to accelerate convergence and iteration time. In particular, we use the heuristic suggested in [7, section 3.4.1] to adjust the ADMM penalty parameter up to 50 times during the optimization procedure; we check the heuristic once every 20 iterations. Sparse Cholesky factorization is used to prefactor the linear system for z before iterations begin; this factorization has to be recomputed any time the penalty parameter changes.

Although ADMM has guaranteed linear convergence in theory, these plots verify that the convergence rates are practical even for large-scale examples, and that they are competitive with highly optimized industrial solvers. Derivative discontinuities in the error plots largely correspond to automatic adjustments of the penalty parameter. Fairly high relative error ($\sim 0.1\%$) is tolerable for our target applications, which typically involve simply plotting the TV profile; from our experiments we can see that ADMM reaches this tolerance level quickly.

Source code including acceleration techniques above is provided in the [GitHub repository](#)² accompanying this paper.

5. Examples and experiments. We evaluate the value of TV profiles through a number of experiments on synthetic shapes as well as American political districts. Our results not only confirm the theoretical properties explored in previous sections but also demonstrate that $I_{\Omega}^{\text{TV}}(t)$ provides a fairly intuitive description of a shape at different length/area scales.

5.1. Synthetic examples. Figure 5 shows the optimized image f for a variety of domains Ω and values of $t \in [0, 1]$. As suggested in section 3.3, for small values of t the images are

²https://github.com/justso1/tv_profile

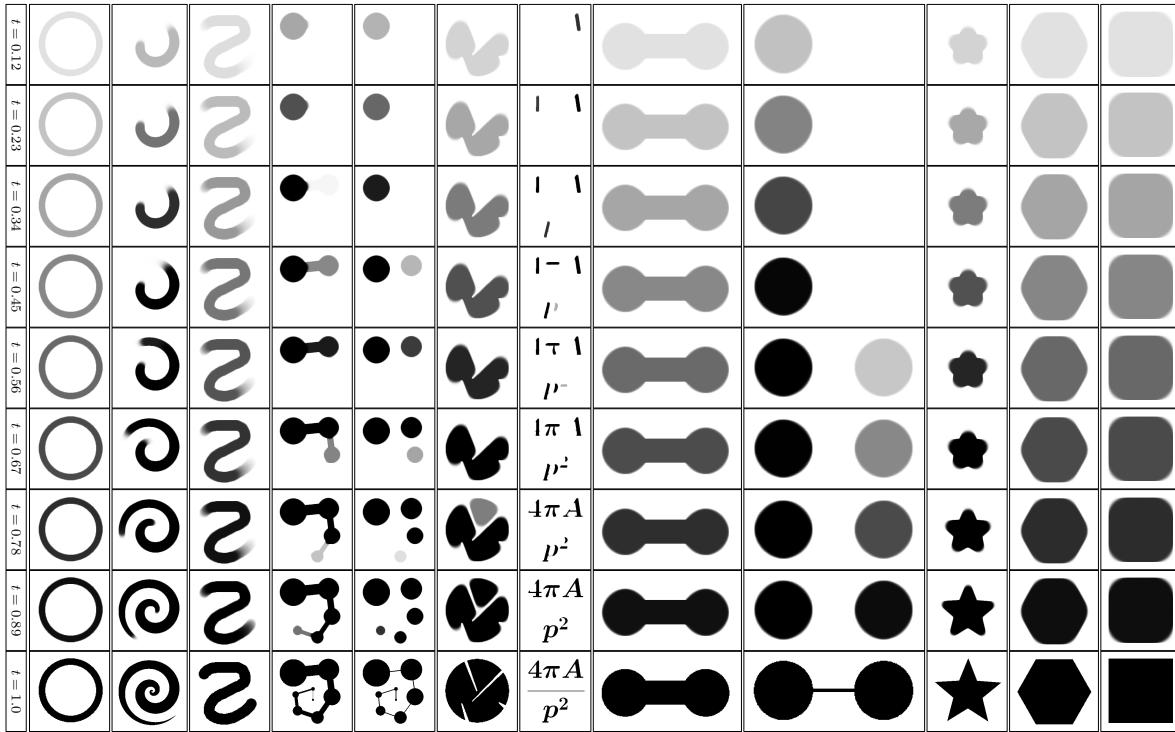


Figure 5. The optimal function f as a function of t for different shapes $\Omega \subseteq \mathbb{R}^2$; here, values of f are scaled from zero (white) to one (black).

simply rescalings of each other, concentrated in a compact subdomain of Ω . For larger values of t , the optimization problem fills in progressively less compact regions within Ω ; at $t = 1.0$ the entire domain is filled.

The corresponding plots of $I_\Omega^{\text{TV}}(t)$ are shown in Figure 6. In these and other plots, we scale the horizontal axis to the range $[0, 1]$ and the vertical axis by the perimeter of a circle whose area is $\text{vol}(\Omega)$. With this scaling, the profile of a circle is the diagonal line with slope 1. As would likely be expected, the text image is the least compact at all t scales. Other more fine-grained information can be obtained by examining these plots, however. For instance, the square and hexagon are considered equally compact at nearly all $t \in [0, 1]$ except near $t \approx 1$. The spiral of circles connected by straight edges has similar compactness values for both edge thicknesses until $t \approx 0.9$, at which point the example with thinner edges distinguishes itself.

For values of t less than 0.75, the dissected circle and the chain of circles connected by increasingly narrow bridges have almost identical TV profiles. Examining Figure 5, we can see that at coarse scales both of the corresponding functions f look like three relatively compact regions connected to each other. At finer scales, however, the dissected circle still appears as three compact connected regions, whereas the three smallest circles in the chain are highly noncompact; this difference leads the two plots to diverge.

We can also observe that the donut shape and the S shape have similar TV profiles at nearly all values of t , with the S being slightly less compact at all resolutions. Again examining

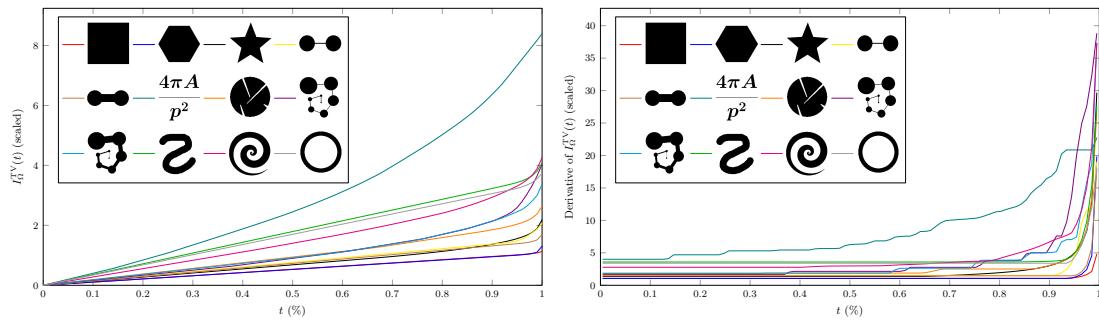


Figure 6. TV profiles (left) and their first derivatives (right) for the shapes in Figure 5.

the images in Figure 5, we can see that the S is highly symmetric and fills in almost uniformly; the donut is truly symmetric and fills in uniformly. Since the figures are of similar width and have similar perimeters, their TV profiles are also similar.

Figure 6 also shows the first derivative of $I_{\Omega}^{\text{TV}}(t)$, computed from the first plot using divided differences. As predicted by Proposition 3.6, all the plots have constant slope starting from $t = 0$. More surprisingly, the plots appear to have several flat regions, suggesting, together with Proposition 3.7 that the set S that appears in the proof of this proposition is discrete. Note that a t for which the derivative of I_{Ω}^{TV} switches from one value to another corresponds to a value for which the isoperimetric profile $I_{\Omega}(t)$ and the TV profile $I_{\Omega}^{\text{TV}}(t)$ coincide.

5.2. Geographic examples. We next evaluate the TV profile on a panel of real congressional districts defined by three different plans for North Carolina’s districts for the U.S. House of Representatives. Following the 2010 Census, North Carolina was apportioned 13 seats, and in 2011 the Republican legislature enacted the districts labeled “2011 districts”³ which were in effect for the 2012 and 2014 congressional elections. After a long court process, this plan was ruled a racial gerrymander by the U.S. District Court for the Middle District of North Carolina in 2016 [26]; the U.S. Supreme Court affirmed the lower court’s decision in 2017 [18], deeming the 2011 plan unconstitutional.

Following these events, the Republican legislature enacted a new map, which took effect before the 2016 elections and is labeled “2016 districts”⁴ in our figures. While more geographically compact, it was again challenged as a racial and partisan gerrymander and was struck down by the same U.S. district court [16]. In 2019, the U.S. Supreme Court overruled the district court, holding that partisan gerrymandering is not federally justiciable, and reinstated the 2016 plan [17].

Additionally, a panel of retired judges simulated a nonpartisan, independent redistricting commission and drew their own proposed plan, henceforth referred to as the *judges’ plan*. An analysis performed by Common Cause North Carolina, using past election data from four election cycles, found that in contrast with the legislature’s enacted plans, which both have ten likely Republican and three likely Democratic districts, the judges’ plan has six likely Republican, four likely Democratic, and three toss-up districts, which more closely matches

³<https://www.census.gov/geo/maps-data/data/tiger-line.html>

⁴<https://www2.ncleg.net/RnR/Redistricting/Main>

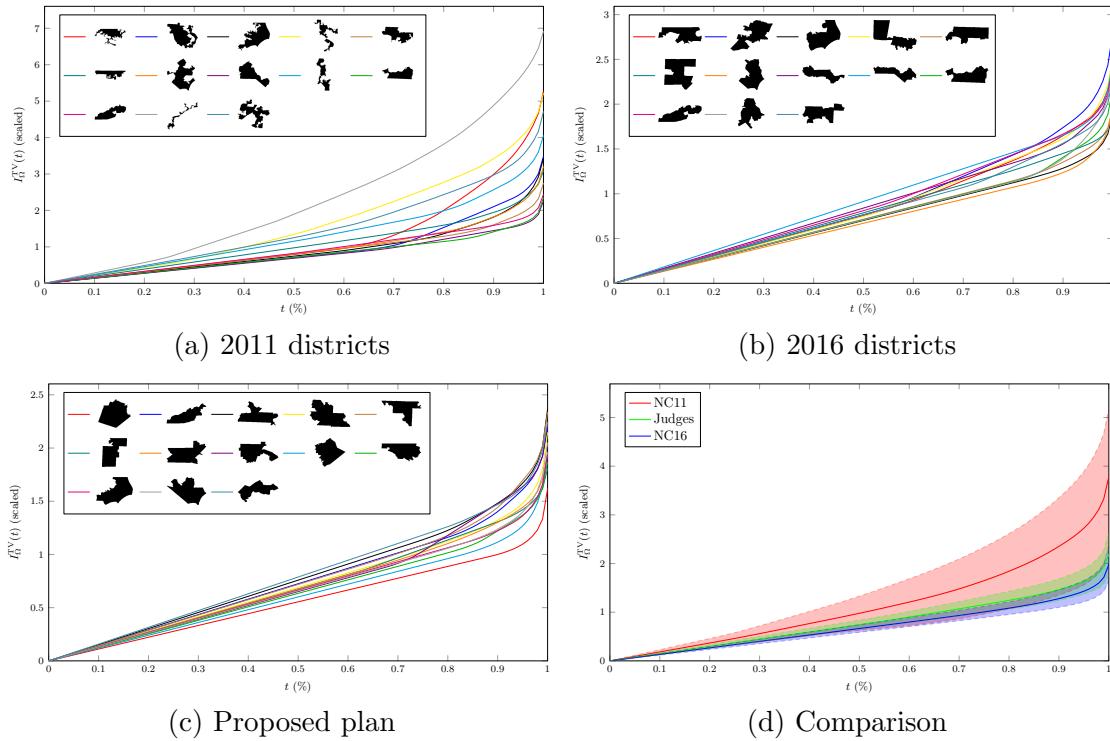


Figure 7. TV profiles for North Carolina districts (a) enacted in 2011 (see Figure 8), (b) enacted in 2016 (see Figure 9), and (c) proposed by a nonpartisan panel of retired judges (see Figure 10). (d) The means and one-standard deviation bands for the TV profiles for the three plans. In (a), (b), and (c) the districts are ordered sequentially from left to right and top to bottom to correspond with the numbering in Figures 8, 9, and 10.

North Carolina's statewide partisan vote shares [46].

We perform our analyses on North Carolina because all three plans are drawn with respect to the 2010 Census population data, reducing the number of confounding factors. Meaningful conclusions about relative compactness are difficult to draw when comparing plans from different states, since factors such as state boundaries, specific laws and rules governing the redistricting process, and number of seats can affect the compactness profile.

For this analysis, we extracted the polygonal boundary of each district, scaled to fit in a 250×250 bounding box; indicator functions of the polygon interiors sampled on this grid were used to approximate $I_O^TV(t)$. Plots of the TV profile for each of the 36 districts are shown in Figures 7(a)–(c). Examining the mean curves and standard deviation bands for the collection of districts in each plan in Figure 7(d), we can see that the 2011 plan appears less compact than the judges' plan, and the judges' plan appears less compact than the 2016 plan. The wide standard deviation band for the 2011 plan is strongly driven by the profile for the twelfth district, which is extremely noncompact. This data reaffirms that noncompactness and gerrymandering are not equivalent: the 2016 map is more compact but less politically representative than the judges' plan.

The curves for the first and ninth districts in the 2011 plan cross around $t = 0.8$, with

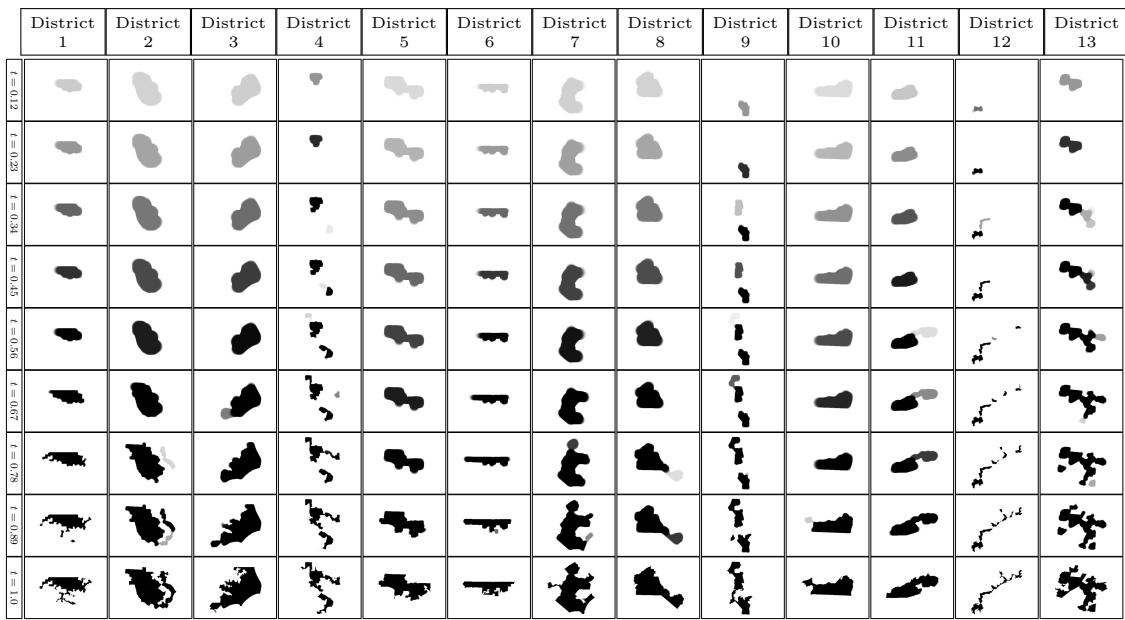


Figure 8. The optimal function f as a function of t for each of the 13 districts in the districting plan enacted in 2011; values of f are scaled from zero (white) to one (black).

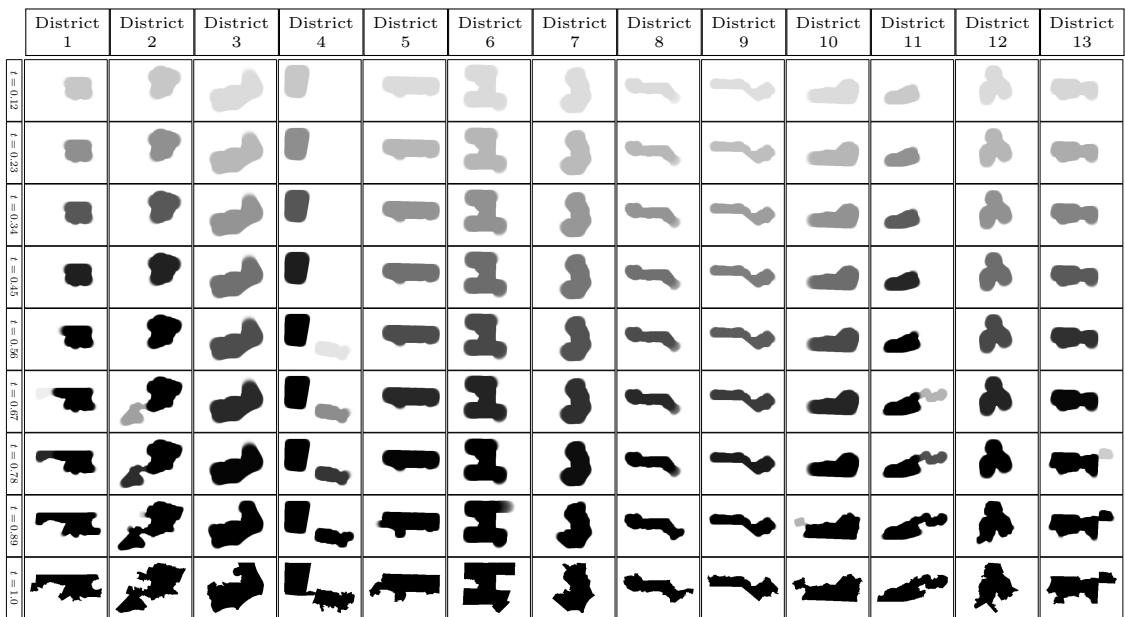


Figure 9. The optimal function f as a function of t for each of the 13 districts in the districting plan enacted in 2016; values of f are scaled from zero (white) to one (black).

the ninth district appearing more compact at larger values of t . This occurs because the ninth district consists of three small compact cores joined with thin connectors, while the first district has a large, fairly compact core on the northern side with several thin, snakelike

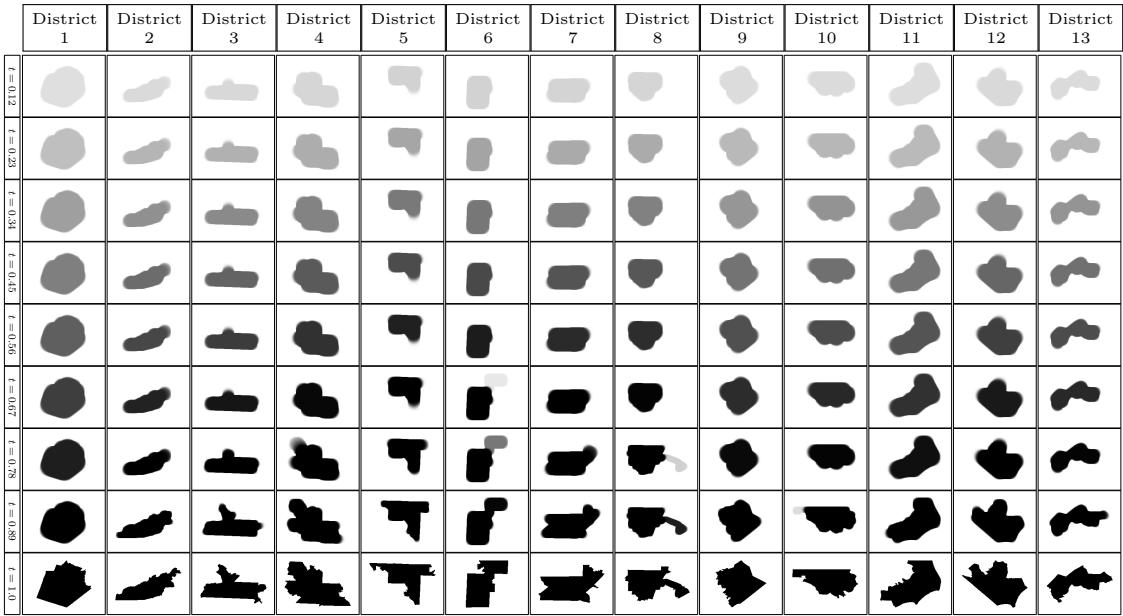


Figure 10. The optimal function f as a function of t for each of the 13 districts in the districting plan proposed by the nonpartisan panel of retired judges; values of f are scaled from zero (white) to one (black).

pieces reaching in various directions. These pieces add a large amount of perimeter and a small amount of area, but the algorithm does not begin filling these appendages until roughly $t = 0.8$; at coarser resolutions, the region is fairly compact.

A similar observation can be made of the eighth district in the judges' plan, which consists of a compact core with a single arm stretching east. At coarse scales, the algorithm fills in this core; given that the other districts are also relatively compact, the curve for the eighth district sits in the middle of this range. Around $t = 0.75$, however, the curve bends upwards and crosses several others. Figure 10 reveals that the algorithm begins filling in this appendage around that time, whereas the algorithm is still filling the core in other districts.

Our analysis above shows the detailed information about compactness encoded in the TV profile. By examining and comparing TV profiles and the corresponding optimized functions f , we can *explain* quantitative scores for evaluating geometric quality.

6. Extensions. We briefly mention some potential extensions of I_Ω^{TV} that may be of interest in different applications. In each case, we show that the extension is a small change of our basic convex optimization problem. Rather than adapting the technique in section 4.2, for simplicity in this section we compute our examples using standard convex cone programming software [35]; our goal is to demonstrate qualitative aspects of these extensions empirically.

6.1. Higher dimensions. The definition (3.2) and discretization (4.2) easily extend to dimensions $n > 2$. This extension allows us to evaluate the compactness of volumes embedded in \mathbb{R}^3 using a cubic lattice discretization of the domain.

Figure 11 shows example shapes computed using this volumetric version of the optimization problem for $I_\Omega^{\text{TV}}(t)$. In this example, we represent $\Omega \subseteq \mathbb{R}^3$ as an indicator on a

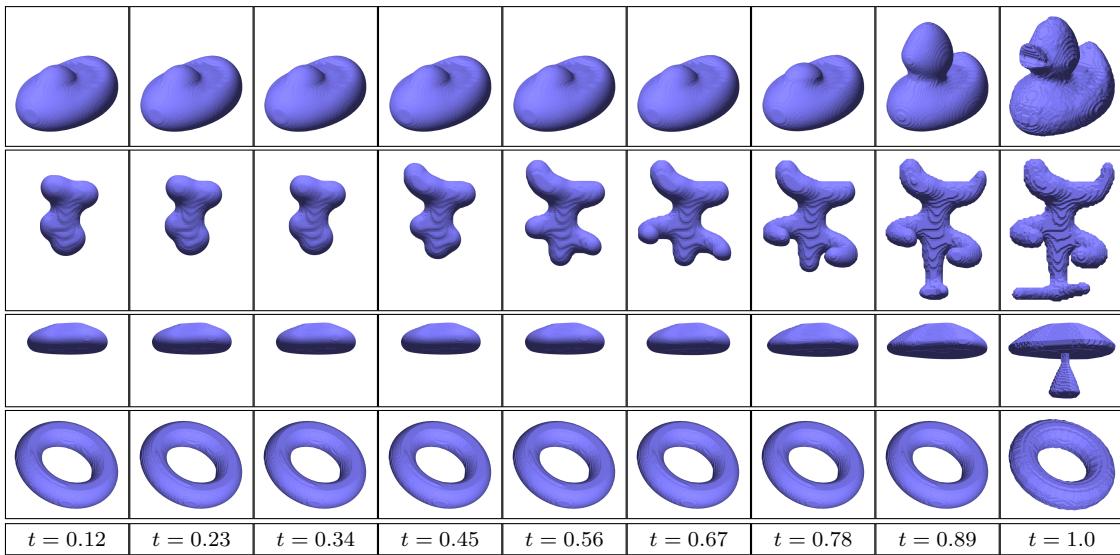


Figure 11. The identical optimization problem for TV profiles can be used for volumetric shapes in \mathbb{R}^3 .

100 × 100 × 100 volume. While gathering enough samples to plot $I_\Omega^{\text{TV}}(t)$ as a function of t is prohibitively expensive, here we show the result of the optimization procedure at sparsely sampled t values; at each t we render the level set of the unknown f at the mean nonzero intensity. Once again, for small t the shape remains fairly constant, as predicted by Proposition 3.6. Simple shapes like the torus remain preserved for most t 's, while more complex shapes like the humanoid start from a nearly convex core when $t \ll 1$ and build up piece by piece.

6.2. Accounting for population density. Our profile (3.2) is defined to be purely geometric and is unaware of any measure on \mathbb{R}^n other than the standard one. Given a distribution $\rho \in \text{Prob}(\mathbb{R}^n)$ —e.g., the population distribution of a state—we can extend our definition to evaluate compactness with respect to ρ :

$$(6.1) \quad I_{\Omega, \rho}^{\text{TV}}(t) := \begin{cases} \min_{f \in L^1(\mathbb{R}^n)} \text{TV}[f] \\ \text{subject to } \int_{\mathbb{R}^n} f(x) d\rho(x) = t, \\ 0 \leq f \leq \mathbb{1}_\Omega. \end{cases}$$

A related formulation is investigated theoretically in [38] for the unrelaxed problem.

Figure 12 shows an experiment using this model. Here, we show the functions $f(x)$ computed on the same district shape, with three different density functions $\rho(x)$. Here we see how the choice of $\rho(x)$ can affect our assessment of compactness; it becomes less expensive to draw circles around densely populated regions contained within the interior of the district, and the boundary of the district becomes less relevant because it is sparsely populated.

6.3. Compactness on a graph. Even though the end result of political redistricting is a collection of geographic districts on a map, redistricting can often be described as a *graph theory* problem. In particular, districts commonly are built out of small geographic subunits,

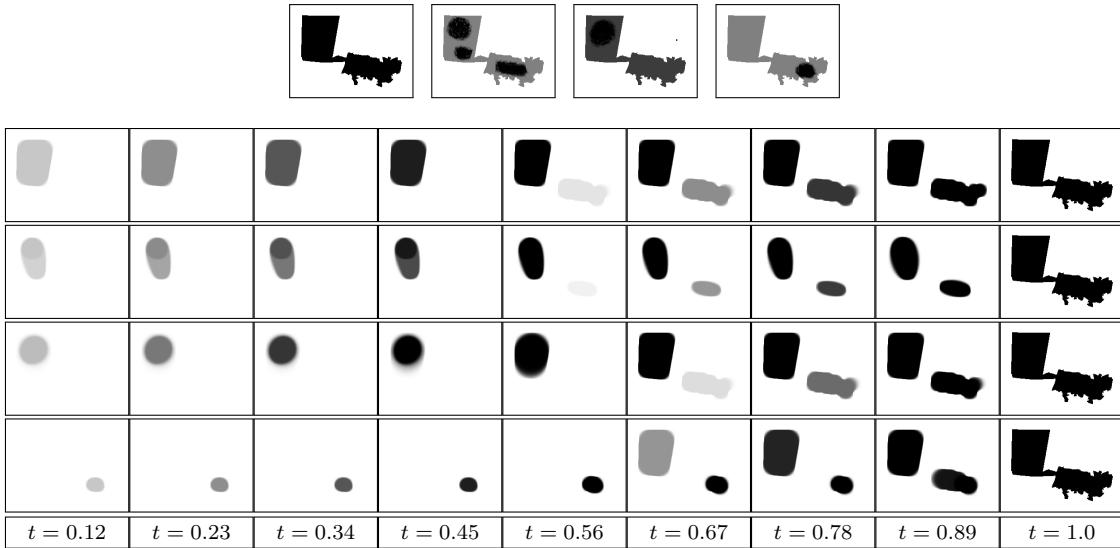


Figure 12. Fictional population distributions in the same district (top, left to right), and the frames of their corresponding weighted TV profiles (bottom, top to bottom).

such as voting tabulation districts (VTDs) or census blocks. Each VTD becomes a vertex in the graph, and two vertices are connected by an edge if and only if their corresponding VTDs are geographically adjacent.⁵ We can then define a graph for each district as the one induced by the vertices corresponding to VTDs in that district. For this reason, it may be of interest to evaluate the compactness of a subset of graph vertices, bypassing embedding into \mathbb{R}^n .

Given a graph $G = (V, E)$, take $V_0 \subseteq V$ to be a subset of vertices representing a potential district. We can imitate (3.2) using graph-based constructions to define a TV profile of V_0 :

$$(6.2) \quad I_{V_0}^{\text{TV}}(t) := \begin{cases} \min_{f \in \mathbb{R}^V} & \sum_{(v,w) \in E} |f(v) - f(w)| \\ \text{subject to} & \sum_{v \in V_0} f(v) = t|V_0|, \\ & f(v) = 0 \ \forall v \notin V_0, \\ & f(v) \in [0, 1] \ \forall v \in V. \end{cases}$$

Note that this definition is not the same as the version of the isoperimetric profile on a graph proposed in [39].

Figure 13 shows the graph-based TV profile for a discretized version of North Carolina's 2011 districting plan, using the same underlying geography as the continuous experiments in section 5.2. To perform the discretization, we use the U.S. Census Bureau's 2010 file of VTDs⁶ and assign each VTD to the district in which it is contained.⁷

⁵ Adjacent here is as according to *rook contiguity*, meaning that two VTDs must share a nontrivial segment of their borders to be considered adjacent. Two VTDs whose borders share only a single point are *not* considered adjacent in our construction.

⁶ <https://www.census.gov/geo/maps-data/data/tiger-line.html>

⁷ Since every VTD is not necessarily contained in a single congressional district, we use *areal interpolation* to assign VTDs to districts, where each is defined to belong to the district with which it shares the most area.

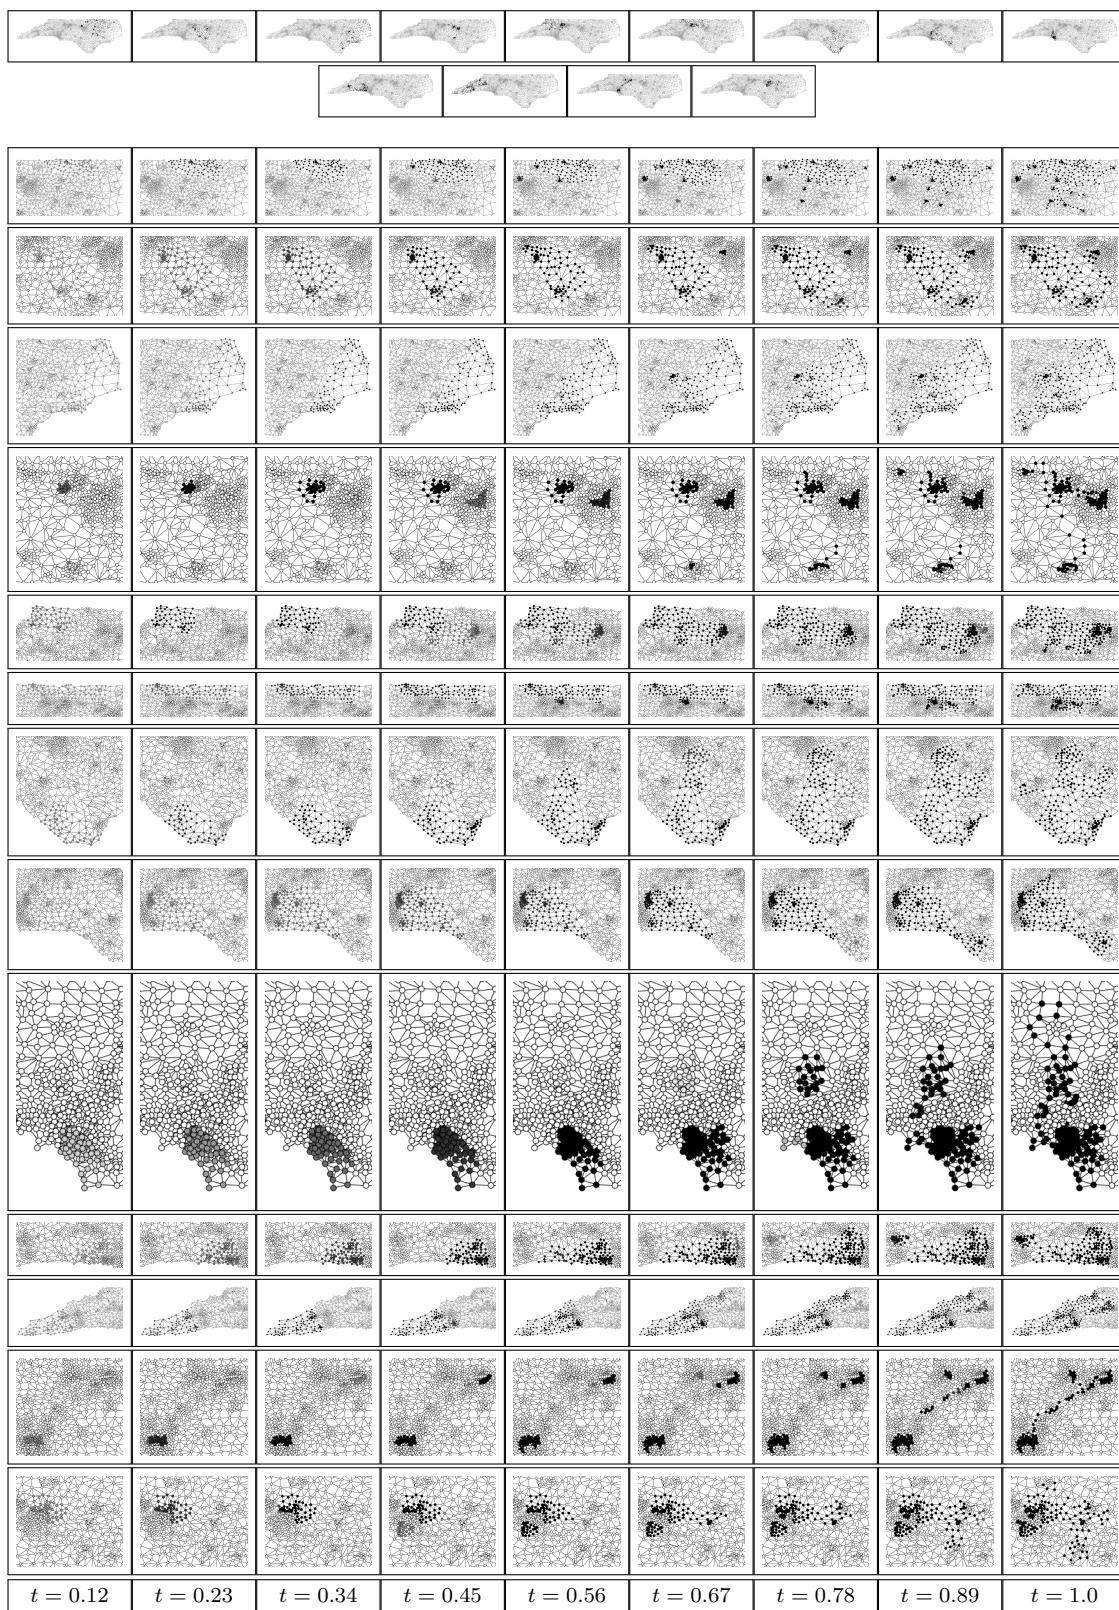


Figure 13. TV profile indicators for different subsets of vertices on a graph of North Carolina voting tabulation districts (VTDs). Images are cropped for each district; the top row shows how the districts are situated in the state. The algorithm is unaware of this embedding and sees only an abstract representation of the graph.

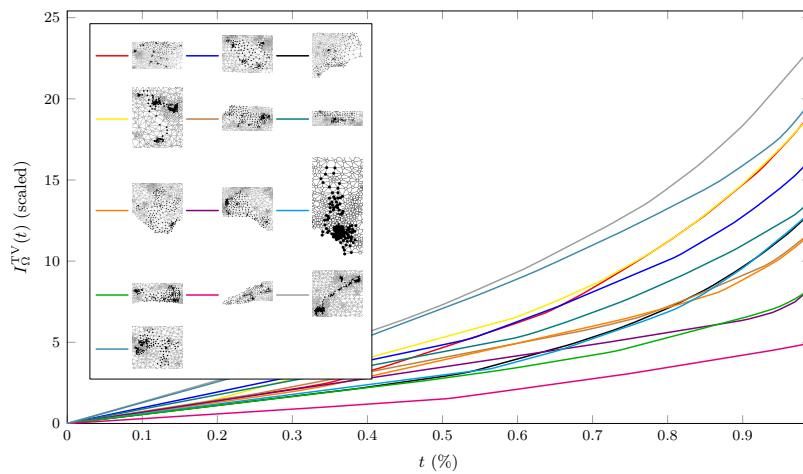


Figure 14. Graph-based TV profiles for the North Carolina districts in Figure 13.

Figure 14 shows the corresponding plots of the TV profile for the discretized version of the 2011 North Carolina districting plan. As in the region-based profiles, the twelfth district is the least compact, and the fourth, tenth, and eleventh are all relatively compact compared to the rest of the districts. Beyond this connection, however, there are many points of difference between the two sets of profiles. One explanation is that areas with higher populations have more VTDs, and hence the graph-based model is in a sense aware of the population distribution. This makes it comparably less expensive to fill in densely populated urban areas, which appear as compact cores under this measure. Conversely, geographically large rural zones have fewer VTDs and therefore appear less compact, as they are more expensive to fill.

To highlight this difference, recall that in the image-based model, most of the tenth district was filled in fairly uniformly, with the small chunk of Asheville in the northwest being filled at large values of t , but under the graph-based algorithm, the bulk of the district is filled in from east to west, beginning in the populous suburbs of Charlotte and spreading through the less heavily populated Western Foothills. Then the portion of the district in Asheville is filled in before these two components are finally connected.

7. Discussion and conclusion. Beyond their interest in the theory of geometry, isoperimetric profiles provide an intuitive, multiscale technique for evaluating the compactness of a shape $\Omega \subseteq \mathbb{R}^n$ that is not completely undermined by the instabilities of the isoperimetric quotient. By considering the entire plot, we gain a fairly complete description of a shape's level of contortion at different length/area scales. Furthermore, our TV-based relaxation of the classical profile admits simpler analysis using convex techniques.

Our initial work in proposing these profiles suggests several avenues for future mathematical, computational, and application-oriented research:

- A key consideration for our long-term target application of political redistricting is to evaluate methods for summarizing these plots and developing useful tests and benchmarks for existing and proposed districting plans. It will also be important to identify the most effec-

tive way to communicate the contents of the TV profile to an end user: nonmathematicians, including politicians and judges, must be able to interpret the results.

- Our formulation currently lacks a formal proof of convergence of our grid-based discretization in the limit of refinement, although this likely is a standard computation adjacent to results in TV-based image denoising.
- Related to convergence, the graph-based model explored experimentally in section 6.3 appears to share many qualities with the TV profile for planar shapes. Exploiting the properties we have developed in the measure-theoretic case may lead to new insight into problems on graphs; ideally the two might be connected formally by examining convergence in the limit of refinement when the graphs are constructed by sampling planar regions.
- Inspired by the Lasserre hierarchy in semidefinite programming [31, 42], we might also ask whether there exist successively tighter convex approximations of the true isoperimetric profile, since our current relaxation is not tight (see Example 3.2); this question is intricately linked with Open Problem 2.1 above.
- Total variation is well-defined for functions on curved manifolds rather than the flat space \mathbb{R}^n ; for instance, we could use our profile to evaluate the compactness of curved segments on a surface, e.g., in the case of geographic data accounting for the curvature and topography of the earth. To compute TV profiles in the presence of curvature, we easily could use a finite element formulation on triangulated surfaces, as suggested in [27].
- Our discussion focuses on total variation as an objective function, but we could attempt to generalize the construction of our profile by considering higher-order measurements popular in mathematical imaging like total generalized variation [10].
- Finally, we could seek uniqueness results about our profile: Is it possible to encounter two shapes Ω with the same TV profiles for all t ?

These open questions aside, given our current analysis and optimization algorithm, the TV profile is already a viable candidate for a nuanced and interpretable multiscale analysis of geometric compactness.

Appendix A. Regularity of the boundary $\partial\Omega$. Throughout the text, the proofs about the properties of the TV profile work as long as

$$(A.1) \quad I_\Omega^{\text{TV}}(t) < +\infty \text{ for } t < \text{vol}(\Omega).$$

Since I_Ω^{TV} is convex (Proposition 3.4), this condition holds as soon as $I_\Omega^{\text{TV}}(t)$ is finite for some t arbitrarily close to $\text{vol}(\Omega)$. We provide two cases where this can be shown to be true.

Proposition A.1. *Assume that $\text{area}(\partial\Omega) < +\infty$. Then (A.1) holds.*

Proof. With this assumption, $I_\Omega^{\text{TV}}(\text{vol}(\Omega)) < +\infty$. Together with the convexity of I_Ω^{TV} and $I_\Omega^{\text{TV}}(0) = 0$, the conclusion clearly holds. ■

In the case where the boundary of Ω is fractal, we can still provide a criterion for (A.1).

Proposition A.2. *Assume that the topological boundary of Ω has zero (n -dimensional) Lebesgue measure. Then (A.1) holds.*

If the boundary of Ω is fractal, in the sense that its d -dimensional Hausdorff measure is finite for some $d \in (n-1, n)$, then we satisfy the assumptions of the proposition above.

Proof. Let D be the distance function to $\mathbb{R}^n \setminus \Omega$. For $\delta < 0$, we set

$$\Sigma_\delta := \{x \in \Omega \text{ such that } D(x) \geq \delta\}.$$

Because of the assumption on the boundary of Ω , $\text{vol}(\Sigma_\delta) \rightarrow \text{vol}(\Omega)$ as $\delta \rightarrow 0$. As the function D is Lipschitz—and hence $\text{TV}(D) < +\infty$ —for a.e. δ we know that $\text{area}(\partial\Sigma_\delta) < +\infty$. As a consequence, using 1_{Σ_δ} as a competitor, $I_\Omega^{\text{TV}}(t) < +\infty$ for a.e. $t \in (0, \text{vol}(\Omega))$. Together with the convexity of I_Ω^{TV} , we see that the conclusion holds. \blacksquare

Appendix B. Duality. Here we carefully justify the duality needed to derive (3.4). The proof relies on the Fenchel–Rockafellar duality theorem [11, Theorem 1.12]. For technical purposes (namely, compactness), we consider $B \subset \mathbb{R}^n$ a closed ball large enough to contain a neighborhood of Ω . All the definitions considered in this paper do not change if we replace \mathbb{R}^n by B .

Let $X = C(B)$ be defined as the space of continuous functions over B valued in \mathbb{R} . Its dual X^* is the set of signed measures over B [44, Theorem 2.14].⁸ Define $F, G : X \rightarrow \mathbb{R} \cup \{+\infty\}$ as

$$F(\eta) := \int_{\Omega} \max(\eta(x), 0) dx$$

and

$$G(\eta) := \begin{cases} -\lambda t & \text{if } \eta = \lambda - \nabla \cdot \phi \text{ with } \lambda \in \mathbb{R}, \phi \in C_c^1(B \rightarrow \mathbb{R}^n) \text{ and } \|\phi\|_{\infty} \leq 1, \\ +\infty & \text{otherwise.} \end{cases}$$

G is well-defined because λ is uniquely defined from η as the mean value of η , if the decomposition $\eta = \lambda - \nabla \cdot \phi$ holds. One can check that F and G are convex. Moreover, at the point $\eta = 0 \in X$, the function F is continuous/finite and G is finite.

Hence, we can apply the Fenchel–Rockafellar duality theorem [11, Theorem 1.12], which states that

$$(B.1) \quad \min_{f \in X} [F(x) + G(x)] = \sup_{\mu \in X^*} [-F^*(\mu) - G^*(-\mu)],$$

where F^*, G^* are the Legendre transforms of F, G . The l.h.s. can be written as

$$\left\{ \begin{array}{ll} \min_{\eta, \lambda, \phi} & -\lambda t + \int_{\Omega} \max(\eta(x), 0) dx \\ \text{subject to} & \|\phi\|_{\infty} \leq 1 \text{ and } \nabla \cdot \phi + \eta = \lambda, \end{array} \right.$$

recovering the r.h.s. of (3.4) up to sign.

On the other hand, we can compute the Legendre transforms of F and G . For $\mu \in X^*$,

$$F^*(\mu) = \sup_{\eta} \int_B \eta \mu - \int_{\Omega} \max(\eta(x), 0) dx.$$

If μ is negative somewhere, taking $\eta \leq 0$ at the same place, one gets larger and larger values for $F^*(\mu)$ by integrating against $s\eta$, where $s \rightarrow +\infty$. On the other hand, if μ is larger than

⁸This would not hold if $B = \mathbb{R}^n$, justifying our introduction of B .

the Lebesgue measure restricted to Ω at some point, taking $\eta \geq 0$ and testing against $s\eta$, $s \rightarrow +\infty$, we reach the conclusion that $F^*(\mu) = +\infty$. In particular, our argument implies that if $F^*(\mu) < +\infty$, then μ has a positive density w.r.t. the restriction of the Lebesgue measure to Ω and that this density is between zero and one. In short,

$$F^*(\mu) = \begin{cases} 0 & \text{if } d\mu(x) = f(x) dx \text{ with } 0 \leq f \leq \mathbb{1}_\Omega \text{ a.e.}, \\ +\infty & \text{otherwise.} \end{cases}$$

As far as G is concerned,

$$\begin{aligned} G^*(\mu) &= \sup_{\eta} \int_B \eta d\mu - G(\eta) \\ &= \sup \left\{ \int_B (\lambda - \nabla \cdot \phi) d\mu + \lambda t : \lambda \in \mathbb{R} \text{ and } \phi \in C_c^1(B \rightarrow \mathbb{R}^n), \|\phi\|_\infty \leq 1 \right\} \\ &= \sup_{\phi} \left\{ \int_B (\nabla \cdot \phi) d\mu : \phi \in C_c^1(B \rightarrow \mathbb{R}^n), \|\phi\|_\infty \leq 1 \right\} + \sup_{\lambda \in \mathbb{R}} \lambda \left(\int_B \mu + t \right). \end{aligned}$$

Provided that μ has a density f w.r.t. the Lebesgue measure (and this is the case if $F^*(\mu) < +\infty$), the first term is none other than $\text{TV}(f)$, and the second is finite if and only if $\int_B f(x) dx = -t$. In short,

$$G^*(f(x) dx) = \begin{cases} \text{TV}(f) & \text{if } \int_B f(x) dx = -t, \\ +\infty & \text{otherwise.} \end{cases}$$

Hence, the r.h.s. of (B.1) reads exactly as $-I_\Omega^{\text{TV}}(t)$, where the latter is defined in (3.2). In conclusion, (B.1) gives the equality between (3.2) and (3.4).

Appendix C. (Generalized) subadditivity for convex function. It is known that convex functions have a tendency of being superadditive. Hence, if they also are subadditive, they must behave like affine functions. This is the subject of the following lemma, which we used in the proof of Proposition 3.7.

Lemma C.1. *Let $F : [0, +\infty) \rightarrow [0, +\infty)$ be a nonnegative convex function and take $r < t \leq s$ such that*

$$(C.1) \quad F(t + s - r) \leq F(t) + F(s) - F(r).$$

Then the function F is affine on $[r, t + s - r]$.

Proof. Let Ψ be the function defined by $\Psi(x) := F(x) + F(s) - F(r) - F(x + s - r)$ (i.e., we consider t as a variable). If Ψ were differentiable, we would easily conclude that it is nonincreasing, but to prove it rigorously we use finite differences. If $r \leq x < y \leq s$,

$$(C.2) \quad \frac{\Psi(y) - \Psi(x)}{y - x} = \frac{F(y) - F(x)}{y - x} - \frac{F(y + s - r) - F(x + s - r)}{y - x} \leq 0,$$

where the last inequality comes from the inequality $x \leq x + s - r$ and the convexity of F .

On the one hand, $\Psi(r) = 0$, and on the other hand, $\Psi(t) \geq 0$ thanks to (C.1). As Ψ is nonincreasing, it is identically 0 on $[r, t]$. Plugging this information back into (C.2) and using

the fact that the finite difference quotients are nondecreasing thanks to convexity of F , we deduce that

$$\frac{F(y) - F(x)}{y - x} = \frac{F(y + s - r) - F(x + s - r)}{y - x} = \text{const.}$$

for any $r \leq x < y \leq t$. Hence the function F is affine on $[r, t]$ and $[s, t + s - r]$ with the same slope, which implies, by convexity of F , that F is affine on $[r, t + s - r]$. \blacksquare

Note that we needed $r < t$ in the proof; otherwise we could not plug such an $x < y$ into (C.2).

Acknowledgments. We thank Erik Demaine, Moon Duchin, Nestor Guillen, Simon Lam, Giorgio Saracco, Filippo Santambrogio, Lily Wang, and the participants of the 2018 Voting Rights Data Institute for their discussions and suggestions provided during the course of this research. We also thank Will Adler for providing areal interpolation code.

REFERENCES

- [1] L. AMBROSIO, N. FUSCO, AND D. PALLARA, *Functions of Bounded Variation and Free Discontinuity Problems*, Clarendon Press Oxford, 2000.
- [2] R. BARNES AND J. SOLOMON, *Gerrymandering and compactness: Implementation flexibility and abuse*, Political Anal., to appear.
- [3] S. BARTELS, *Total variation minimization with finite elements: Convergence and iterative solution*, SIAM J. Numer. Anal., 50 (2012), pp. 1162–1180, <https://doi.org/10.1137/11083277X>.
- [4] A. BESICOVITCH, *A variant of a classical isoperimetric problem*, Quart. J. Math. Oxford Ser., 20 (1949), pp. 84–94.
- [5] A. BESICOVITCH, *Variants of a classical isoperimetric problem*, Quart. J. Math. Oxford Ser. (2), 3 (1952), pp. 42–49.
- [6] J. BORWEIN AND A. S. LEWIS, *Convex Analysis and Nonlinear Optimization: Theory and Examples*, Springer Science & Business Media, 2010.
- [7] S. BOYD, N. PARikh, E. CHU, B. PELEATO, AND J. ECKSTEIN, *Distributed optimization and statistical learning via the alternating direction method of multipliers*, Found. Trends Mach. Learning, 3 (2011), pp. 1–122.
- [8] S. BOYD AND L. VANDENBERGHE, *Convex Optimization*, Cambridge University Press, 2004.
- [9] A. BRAIDES, *Approximation of Free-Discontinuity Problems*, Lecture Notes in Math. 1694, Springer Science & Business Media, 1998.
- [10] K. BREDIES, K. KUNISCH, AND T. POCK, *Total generalized variation*, SIAM J. Imaging Sci., 3 (2010), pp. 492–526, <https://doi.org/10.1137/090769521>.
- [11] H. BREZIS, *Functional Analysis, Sobolev Spaces and Partial Differential Equations*, Springer Science & Business Media, 2010.
- [12] G. CARLIER, M. COMTE, AND G. PEYRÉ, *Approximation of maximal Cheeger sets by projection*, M2AN Math. Model. Numer. Anal., 43 (2009), pp. 139–150.
- [13] A. CHAMBOLLE, V. CASELLES, D. CREMERS, M. NOVAGA, AND T. POCK, *An introduction to total variation for image analysis*, in Theoretical Foundations and Numerical Methods for Sparse Recovery, Radon Ser. Comput. Appl. Math. 9, Walter de Gruyter, Berlin, 2010, pp. 263–340.
- [14] A. CHAMBOLLE, V. DUVAL, G. PEYRÉ, AND C. POON, *Geometric properties of solutions to the total variation denoising problem*, Inverse Problems, 33 (2016), 015002.
- [15] A. CHAMBOLLE AND T. POCK, *An introduction to continuous optimization for imaging*, Acta Numer., 25 (2016), pp. 161–319, <https://doi.org/10.1017/S096249291600009X>.
- [16] *Common Cause v. Rucho*, United States District Court, Middle District of North Carolina, 2017.
- [17] *Common Cause v. Rucho*, Supreme Court of the United States, 2019.
- [18] *Cooper v. Harris*, Supreme Court of the United States, 2017.

- [19] H. T. CROFT, K. FALCONER, AND R. K. GUY, *Unsolved Problems in Geometry: Unsolved Problems in Intuitive Mathematics*, Vol. 2, Springer Science & Business Media, 2012.
- [20] M. DUCHIN AND B. E. TENNER, *Discrete Geometry for Electoral Geography*, preprint, <https://arxiv.org/abs/1808.05860>, 2018.
- [21] K. EHRENBURG, *Studien zur messung der horizontalen gliederung von erdräumen*, Verhandlungen der Physikalisch-medizinische Gesellschaft zu Würzburg, 25–26 (1892), pp. 41–92.
- [22] W. H. FLEMING AND R. RISHEL, *An integral formula for total gradient variation*, Arch. Math. (Basel), 11 (1960), pp. 218–222.
- [23] E. GONZALEZ, U. MASSARI, AND I. TAMANINI, *On the regularity of boundaries of sets minimizing perimeter with a volume constraint*, Indiana Univ. Math. J., 32 (1983), pp. 25–37.
- [24] M. GRANT AND S. BOYD, *Graph implementations for nonsmooth convex programs*, in Recent Advances in Learning and Control, V. Blondel, S. Boyd, and H. Kimura, eds., Lect. Notes Control Inf. Sci. 371, Springer-Verlag, 2008, pp. 95–110, http://stanford.edu/~boyd/graph_dcp.html.
- [25] M. GRANT AND S. BOYD, *CVX: Matlab Software for Disciplined Convex Programming*, Version 2.1, <http://cvxr.com/cvx>, 2014.
- [26] *Harris v. McCrory*, United States District Court, Middle District of North Carolina, 2016.
- [27] M. HERRMANN, R. HERZOG, S. SCHMIDT, J. VIDAL-NÚÑEZ, AND G. WACHSMUTH, *Discrete Total Variation with Finite Elements and Applications to Imaging*, preprint, <https://arxiv.org/abs/1804.07477>, 2018.
- [28] A. KAUFMAN, G. KING, AND M. KOMISARCHIK, *How to measure legislative district compactness if you only know it when you see it*, Amer. J. Polit. Sci., to appear, <http://j.mp/Compactness>.
- [29] B. KAWOHL AND T. LACHAND-ROBERT, *Characterization of Cheeger sets for convex subsets of the plane*, Pacific J. Math., 225 (2006), pp. 103–118.
- [30] M.-J. LAI, B. LUCIER, AND J. WANG, *The convergence of a central-difference discretization of Rudin-Osher-Fatemi model for image denoising*, in International Conference on Scale Space and Variational Methods in Computer Vision, Springer, 2009, pp. 514–526.
- [31] J. B. LASSERRE, *Convexity in semialgebraic geometry and polynomial optimization*, SIAM J. Optim., 19 (2009), pp. 1995–2014, <https://doi.org/10.1137/080728214>.
- [32] G. P. LEONARDI, R. NEUMAYER, AND G. SARACCO, *The Cheeger constant of a Jordan domain without necks*, Calc. Var. Partial Differential Equations, 56 (2017), 164.
- [33] G. P. LEONARDI AND A. PRATELLI, *On the Cheeger sets in strips and non-convex domains*, Calc. Var. Partial Differential Equations, 55 (2016), 15.
- [34] G. P. LEONARDI AND G. SARACCO, *Minimizers of the prescribed curvature functional in a Jordan domain with no necks*, in preparation.
- [35] MOSEK APS, *The MOSEK Optimization Toolbox for MATLAB*, 2017, <http://docs.mosek.com/8.1/toolbox/index.html>.
- [36] E. PARINI, *An introduction to the Cheeger problem*, Surv. Math. Appl., 6 (2011), pp. 9–21.
- [37] D. D. POLSBY AND R. D. POPPER, *The third criterion: Compactness as a procedural safeguard against partisan gerrymandering*, Yale Law Policy Rev., 9 (1991), pp. 301–353.
- [38] A. PRATELLI AND G. SARACCO, *On the isoperimetric problem with double density*, Nonlinear Anal., 177 (2018), pp. 733–752.
- [39] P. RAGHAVENDRA, D. STEURER, AND P. TETALI, *Approximations for the isoperimetric and spectral profile of graphs and related parameters*, in Proceedings of the 2010 ACM International Symposium on Theory of Computing (STOC’10), ACM, 2010, pp. 631–640.
- [40] R. T. ROCKAFELLAR, *Convex Analysis*, Princeton University Press, 1997.
- [41] A. ROS, *The isoperimetric problem*, in Global Theory of Minimal Surfaces, Vol. 2, AMS, 2001, pp. 175–209.
- [42] T. ROTHVOSS, *The Lasserre hierarchy in approximation algorithms*, Lecture Notes for the MAPSP, 2013.
- [43] L. I. RUDIN, S. OSHER, AND E. FATEMI, *Nonlinear total variation based noise removal algorithms*, Phys. D, 60 (1992), pp. 259–268.
- [44] W. RUDIN, *Real and Complex Analysis*, 3rd ed., McGraw-Hill, 1987.
- [45] E. STREDULINSKY AND W. P. ZIEMER, *Area minimizing sets subject to a volume constraint in a convex set*, J. Geom. Anal., 7 (1997), pp. 653–677.
- [46] B. WARNER, *Retired Judges Unveil New NC Congressional Map*, 2016, <https://www.commoncause.org/north-carolina/press-release/retired-judges-unveil-new/>.