

Spectral Total-Variation Processing of Shapes - Theory and Applications

JONATHAN BROKMAN, Technion, Israel Institute of Technology

MARTIN BURGER, Helmholtz Imaging, Deutsches Elektronen-Synchrotron DESY, Germany and Fachbereich Mathematik, Universität Hamburg, Germany

GUY GILBOA, Technion, Israel Institute of Technology

We present an analysis of total-variation (TV) on non-Euclidean parameterized surfaces, a natural representation of the shapes used in 3D graphics. Our work explains recent experimental findings in shape spectral TV [Fumero et al. 2020] and adaptive anisotropic spectral TV [Biton and Gilboa 2022]. A new way to generalize set convexity from the plane to surfaces is derived by characterizing the TV eigenfunctions on surfaces. Relationships between TV, area, eigenvalue, eigenfunctions and their discontinuities are discovered. Further, we expand the shape spectral TV toolkit to include versatile zero-homogeneous flows demonstrated through smoothing and exaggerating filters. Last but not least, we propose the first TV-based method for shape deformation, characterized by deformations along geometrical bottlenecks. We show these bottlenecks to be aligned with eigenfunction discontinuities. This research advances the field of spectral TV on surfaces and its application in 3D graphics, offering new perspectives for shape filtering and deformation.

CCS Concepts: • **Mathematics of computing**;

Additional Key Words and Phrases: geometry processing, total-variation, 3-Laplacian, nonlinear, non-Euclidean, nonlinear spectral processing

1 INTRODUCTION

Spectral geometry processing is a widely used technique in computer graphics. It involves breaking down shapes and their functions into spectral components, harnessing the multiscale nature and the well studied properties of eigenfunctions [Aflalo et al. 2013; Bracha et al. 2020; Sorkine et al. 2004; Taubin 1995; Wetzler et al. 2013]. Spectral processing is useful for various computer graphics applications, such as mesh filtering, surface decoding, segmentation (spectral partitioning), shape deformation and general analysis. For a comprehensive exploration of this approach, including a survey and overview, please refer to [Cammaraana and Patané 2021], and the references therein.

Generally, the spectral approach is based on linear algebra and harmonic analysis. It can be viewed as a generalization of Fourier basis to Riemannian manifolds and to graphs. In recent years, there has been a growing branch studying nonlinear spectral decompositions, see e.g. [Gilboa 2018]. In this setting, a signal is decomposed into spectral components related to nonlinear eigenfunctions of the form $\lambda u \in \partial J(u)$, where ∂J is the subdifferential of a convex functional J . It can be readily seen that when J is the Dirichlet energy, $J(u) = \|\nabla u\|_2^2$, we have $\partial J(u) = \nabla J(u) = -\Delta u$, where Δ denotes the Laplacian. Thus, the eigen-decomposition reduces back to a linear Fourier-type analysis. In this sense, it is a generalization of linear spectral methods. Most of the research was dedicated to absolutely one-homogeneous functionals [Bungert et al. 2019; Burger




Fig. 1. Spectral TV filtering of a shape, demonstrating the application of different filters to isolate TV eigenfunctions and reveal geometric details. Left: The original shape with an all-pass filter, displaying the complete spectrum. Middle: A low-pass filter uncovers the shape's fundamental structure. From a theoretical point of view - it isolates eigenfunctions as lower-frequency components. Right: A high-pass filter brings forth the geometrical details on-top the foundational structure. The geometry is partitioned into three distinct subsets based on eigenfunction discontinuities (details in Fig. 6). This delineates the original shape by its bottlenecks, which are also observed in the boundaries of each of the three nested, concentric surfaces of the right panel. These bottleneck boundaries are a direct manifestation of the eigenfunction properties, which our manuscript investigates through novel theoretical and empirical analysis.

et al. 2016; Gilboa 2014a] such as total variation. These new signal analysis methods enabled crisp, edge-preserving decompositions and representations. However, the theory and applications were restricted mostly to images and signals in Euclidean domains. In this study we develop a theoretical setting for nonlinear spectral processing on surfaces, demonstrating the potential benefits of this approach to computer graphics.

Total Variation (TV) is a popular regularization functional, well-known for its edge-preserving properties [Chambolle 2004]. For a smooth function $u : \Omega \rightarrow \mathbb{R}$, the TV functional is

$$TV(u) = \int_{\Omega} |\nabla u(x)| dx. \quad (1)$$

Total variation has found extensive applications in various domains and tasks over the past three decades. Most notably, this approach was extensively used in image processing, for denoising, deconvolution, texture modeling, super-resolution, segmentation and more [Aujol et al. 2006; Burger and Osher 2013a; Chambolle and Pock

Authors' addresses: Jonathan Brokman, jonathanbrok@gmail.com, Technion, Technion City, Haifa, Israel Institute of Technology, 32000; Martin Burger, martin.burger@desy.de, Helmholtz Imaging, Deutsches Elektronen-Synchrotron DESY, Notkestr. 85, Hamburg, Germany, 22607 and Fachbereich Mathematik, Universität Hamburg, Bundesstr. 55, Hamburg, Germany, 20146; Guy Gilboa, guy.gilboa@ee.technion.ac.il, Technion, Technion City, Haifa, Israel Institute of Technology, 32000.




Fig. 2. "Bust of Queen Nefertiti"¹: Nonlinear spectral filtering applied using one of our proposed methods, presented in Sec. 7.5

2011; Leng et al. 2021; Pascal et al. 2021; Rudin et al. 1992; Wei et al. 2019; Zhang et al. 2022].

The exploration of Total Variation extends beyond images and encompasses shapes, point clouds, and graphs [Dinesh et al. 2020; Duan et al. 2019; Sawant and Prabukumar 2020]. Consequently, it is natural to harness the TV functional for its nonlinear spectral properties in fields like spectral geometry. Such an approach was initiated by [Fumero et al. 2020], which introduced Spectral Total Variation [Gilboa 2013, 2014b] for geometry analysis. Nonetheless, transitioning from images to shapes presents new challenges.

In the domain of image processing, regularization predominantly revolves around manipulating a function u defined within a Euclidean domain. However, when it comes to geometry processing, the usage of u as a representation of the processed shape introduces a significant distinction, as the assumption of a Euclidean metric is typically not valid. For Spectral TV, this distinction sparks unique and unexplored considerations.

The primary focus of this paper centers around the theory of TV eigenfunctions and Spectral TV within the realm of 3D geometry. Our main contributions are:

- A new way to generalize the definition of convex sets from planar domains to surfaces. This generalization arises as a property of total variation (TV) eigenfunctions on non-Euclidean surfaces, addressing theoretical inquiries of recent years [Fumero et al. 2020], [Biton and Gilboa 2022].
- Further quantitative relationships between eigenvalue, area, and the total variation of the eigenfunctions are derived.
- Numerical demonstrations of our theoretical findings on the eigenfunction properties are provided, Figs. 4, 5, 6.
- Spectral TV is extended to include novel zero-homogeneous flows, utilized for spectral filtering of shapes, where each flow induces distinct filtering capabilities - while being multi-scale as expected from spectral methods. See Figs. 9, 11, 21.

¹"Bust of Queen Nefertiti". Ägyptisches Museum und Papyrussammlung. The original sculpture was created in 1345 BC by Thutmose. Scanned by Nora Al-Badri and Jan Nikolai Nelles. Available under a [Creative Commons license](#).

- We present the first TV based solution to the Shape Deformation task, as described in the tutorial by [Sorkine and Botsch 2009]. This approach typically results in the concentration of deformation around geometrical bottlenecks, see Figs. 15, 19.
- The discovered properties of eigenfunctions are observed in both our shape filtering and shape deformation results, enhancing the understanding of these methods, notably in Figs. 6, 12, 13, 14.

Through these contributions, we not only derive new generalizations of key theoretical concepts from Euclidean planes to non-Euclidean surfaces but also showcase how these theoretical advancements can be applied to shape processing tasks, offering new insights as well as methodologies. For 3D model details, see footnote under references. Our results are reproducible via our [official implementation](#).

2 PREVIOUS WORK

Shape filtering has a long and rich history. The pivotal work of [Taubin 1995] proposed to utilize the shape-induced Laplacian eigenfunctions as a basis for shape filtering, in an analogue manner to classical signal processing techniques. A transform is computed by projecting the shape onto the basis, where filtering is obtained by weighted reconstruction via this basis. Many variations of this method were utilized for different tasks (e.g. [Sorkine et al. 2004]). Over time, the Laplace-Beltrami became the standard Laplacian of choice, for spectral applications, and in general [Wetzler et al. 2013]². For non-rigid shape processing, spectral representations of a scale-invariant version of the Laplace-Beltrami was introduced in [Aflalo et al. 2013], and later combined with deep learning for shape correspondence [Bracha et al. 2020], providing a non-rigid spectral analogue of the popular spatial-domain approach of [Litany et al. 2017].

While Total Variation (TV) is traditionally used for image processing (see [Burger and Osher 2013b; Chambolle et al. 2010] for theory and applications), it was used in computer graphics as well, for surface denoising, reconstruction, segmentation, and super-resolution [Dinesh et al. 2019, 2020; Kerautret and Lachaud 2020; Liu et al. 2017; Zhang and Wang 2020; Zhang et al. 2015; Zhong et al. 2018]. Only recently, in 2020, a TV-based shape filtering approach was proposed [Fumero et al. 2020], advocating the application of spectral TV to the normals of the shape.

TV and shape smoothing flows have a substantial background, spanning a considerable duration. The authors of [Kimmel et al. 1998] established a close relation between TV image smoothing and Laplace-Beltrami shape smoothing. These techniques were based on discrete diffusion processes. Among a plethora of Laplacian-based diffusion processes we note two fundamental flows: the Mean Curvature Flow (MCF), and the Willmore flow, which gave rise to a new problem of singularities [Blatt 2009; Desbrun et al. 1999; Huisken 1990]. More stable variations of them were proposed throughout the years, for instance [Crane et al. 2013] and [Kazhdan

²Such an adaptation of [Taubin 1995] can be found e.g. in [Vallet and Lévy 2008], which proposed a computationally efficient shape filtering, and demonstrated some core filtering capabilities: Shape exaggeration, detail enhancement, shape smoothing and regularization.

et al. 2012]. In [Elmoataz et al. 2008] a generalization of the Laplacian-based approaches was proposed. Using graph-oriented operators they introduced p -Laplace flows for mesh fairing, including the 1-Laplace flow, which is a TV-flow, the gradient flow minimizing the TV energy.

Various methods have been proposed to minimize numerically TV and a L^2 square fidelity term ($J_{TV}(u) + \lambda \|f - u\|_2^2$, often referred to as the ROF model [Rudin et al. 1992]), see for instance [Chambolle 2004; Chambolle and Pock 2011]. For implementing a TV-flow, one should note that the term $-\operatorname{div}(\frac{\nabla u}{|\nabla u|})$ is valid for the gradient estimation of the energy only for non-vanishing gradients of u . Two common approaches are adopted to overcome this. One is to use a regularized TV model $J_{TV-\varepsilon} := \int_{\Omega} \sqrt{|\nabla u(x)|^2 + \varepsilon^2} dx$. In this case, one does not obtain singularities at zero gradients and a straightforward explicit method can be used. Drawbacks of this approach are that an additional ε parameter is introduced and that the flow is smoother (somewhat less edge preserving). Moreover, the evolution time step is highly restrictive, proportional to ε . Alternatively, one can approximate the gradient descent process as a series of proximal TV minimizations and use non-smooth solvers to obtain fully-edge preserving solutions, e.g. by [Chambolle 2004] or [Chambolle and Pock 2011]. This yields more faithful results. These methods can be regarded as semi-implicit methods, which are unconditionally stable (with respect to the time step parameter, as opposed to explicit methods).

In our work we adapt two such processes for geometry processing: The first is the shape flow introduced in [Kazhdan et al. 2012], which proposed a gradient descent semi-implicit flow (see [Parikh et al. 2014] for semi-implicit approaches). It is adapted by changing their proposed gradient direction to conform with our setting. The second is the iterative re-weighted L1 minimization scheme, introduced in [Bronstein et al. 2016], which is adapted by using a vectorial version of it, and again, selecting our operator of choice to replace their gradient direction. The shape deformation problem is classically solved as a constrained energy-minimization problem, for instance [Botsch and Sorkine 2007; Sorkine and Alexa 2007], which we also explore in our work.

Another approach which relates to our method is the family of shape processing techniques which process the displacement fields over a smooth version of the surface. This is mainly used for shape smoothing, exaggeration and also for detail transfer, for instance [Cignoni et al. 2005; Digne 2012; Sorkine and Botsch 2009]. Recently, [Yifan et al. 2021] leveraged two separate neural networks for the over-smoothed shape and for the displacement.

Our most significant contribution is to the theory of spectral TV. Let us recap significant landmarks of this domain for the Euclidean setting. Spectral TV was introduced in [Gilboa 2013, 2014b], facilitating nonlinear edge-preserving image filtering. Essentially, the idea is to decompose a signal into nonlinear spectral elements related to eigenfunctions of the total-variation subdifferential. These nonlinear eigenfunctions raise important theoretical aspect - as their properties enable an understanding of the behaviour of both TV regularization, and Spectral Total Variation processing. The method is based on evolving gradient descent with respect to the TV functional, also known as the TV-flow [Andreu et al. 2001a].

The spectral elements decay linearly in this flow. Different decay rates correspond to different scales, where in the case of a single eigenfunction the rate is exactly the eigenvalue. Theoretical underpinning was performed for the spatially discrete case in [Burger et al. 2016], which also extended the concepts of spectral TV to decompositions based on general convex absolutely one-homogeneous functionals. Decompositions based on minimizations with the Euclidean norm, as well as with inverse-scale-space flows [Burger et al. 2006] were also proposed. The space-continuous setting was later analyzed in [Bungert et al. 2019]. For the one-dimensional TV setting, it was shown that the spectral elements are orthogonal to each other. Various applications were suggested for image enhancement, manipulation and fusion [Benning et al. 2017; Hait and Gilboa 2019]. A common thread related to gradient flows of one-homogeneous functionals is that they are based on zero-homogeneous operators. Note that other homogeneities are also possible, see for instance [Bungert and Burger 2020; Cohen and Gilboa 2020].

While in non-Euclidean domains the theory of TV eigenfunctions and spectral TV is not as developed as in the Euclidean case, other aspects of TV on non-Euclidean domains are highly researched. A TV framework on manifolds was defined in [Ben-Artzi and LeFloch 2007]. This research was in the context of nonlinear hyperbolic conservation laws on Riemannian manifolds. The authors prove bounds and stability of the minimizing flow. Another example is anisotropic TV, which is typically applied for image processing (see [Grasmair and Lenzen 2010]), and can be analyzed as a functional on a non-Euclidean domain - as shown in [Biton and Gilboa 2022], which conducted experiments with spectral TV as well. In [Fumero et al. 2020] experiments with geodesically convex sets³ were conducted. In their experiments, geodesically convex sets exhibited a linear decay throughout the minimizing flow. Theoretical explanations of this phenomenon were left as opened research questions. In the following we lay theoretical foundations for these experimental results, validated by our proposed spectral nonlinear and non-Euclidean framework.

3 PRELIMINARIES

In this work, the processed shape is assumed to be a 2D manifold $M \subset \mathbb{R}^3$. We assume M is smooth, with a smooth parameterization⁴ $S(\omega_1, \omega_2) = (x(\omega_1, \omega_2), y(\omega_1, \omega_2), z(\omega_1, \omega_2))$, $\omega_1, \omega_2 \in \Omega$, i.e.

$$S : \Omega \subset \mathbb{R}^2 \rightarrow M. \quad (2)$$

Namely, S is differentiable and invertible. In the following we outline important well-known properties of this setting, which we use in our work. A celebrated resource for these properties is [Do Carmo 2016].

Let $f : M \rightarrow \mathbb{R}$, and $u = f \circ S(\omega_1, \omega_2)$ i.e. $u : \Omega \rightarrow \mathbb{R}$. Let $T_q M$ be the plane tangent to M at point $S(\omega_1, \omega_2) = q \in M$. It can be shown that $\frac{\partial S}{\partial \omega_1}, \frac{\partial S}{\partial \omega_2} \in \mathbb{R}^3$, denoted $S_{\omega_1}, S_{\omega_2}$, span $T_q M$ at point q , assuming they are linearly independent. A field F on M assigns

³Geodesically convex sets are subsets of a manifold in which any two points have the geodesics between them contained in the set. Sometimes, it is further assumed that there is a unique geodesic between two such points.

⁴ M locally behaves like \mathbb{R}^2 . See for instance smooth manifold as "coordinate system" defined in Thm. 5-2 of [Spivak 2018]. See also local diffeomorphism for instance in [Lee 2013].

each q with a vector in T_q , i.e. $F : M \rightarrow \cup(T_q)_{q \in M}$. Let $[t_1, t_2] \subset \mathbb{R}$ a connected interval. We denote a differentiable parametric curve $(\omega_1(t), \omega_2(t))$, $t \in [t_1, t_2]$ by $\tilde{\gamma}(t) : [t_1, t_2] \rightarrow \Omega$, and its mapped curve by $\gamma = S(\tilde{\gamma}(t)) \subset M$, i.e. γ is a differentiable parametric curve which maps $\gamma : [t_1, t_2] \rightarrow M$. We write $\frac{d\gamma}{dt}$ as γ_t , often interpreted as the velocity at time t . Using the chain rule, γ_t can be calculated as $\gamma_t = S_{\omega_1} \frac{\partial \omega_1}{\partial t} + S_{\omega_2} \frac{\partial \omega_2}{\partial t} = J \tilde{\gamma}_t$ where $J(\omega_1, \omega_2)$ is the Jacobian matrix with columns $(S_{\omega_1}, S_{\omega_2})$. For a vector \tilde{a} originating from a point ω_1, ω_2 , similarly to the velocity vectors $\tilde{\gamma}_t$, \tilde{a} is mapped to M as

$$a(q) = J(\omega_1, \omega_2) \tilde{a}, \quad (3)$$

where $q = S(\omega_1, \omega_2)$. Note that $a(q) \in T_q$. The induced inner product on Ω is $\langle \tilde{a}, \tilde{b} \rangle_g = (J\tilde{a})^T (J\tilde{b}) = \tilde{a}^T (J^T J) \tilde{b}$. Hence, when M is parameterized using (Ω, S) , it is equipped with the metric $g = J^T J$:

$$g(\omega_1, \omega_2) = \begin{pmatrix} S_{\omega_1}^T S_{\omega_1} & S_{\omega_1}^T S_{\omega_2} \\ S_{\omega_2}^T S_{\omega_1} & S_{\omega_2}^T S_{\omega_2} \end{pmatrix}. \quad (4)$$

$S_{\omega_1}, S_{\omega_2}$ are assumed to be linearly independent, thus g is positive definite (and invertible). The vector magnitude can now be calculated using \tilde{a}, g as $\|a\|_2 = \|\tilde{a}\|_g := \|J\tilde{a}\|_2 = \sqrt{\langle \tilde{a}, \tilde{a} \rangle_g}$. As a consequence, the length of γ is $L(\gamma) = \int_c^d \|\gamma_t\|_2 dt = \int_c^d \|\tilde{\gamma}_t\|_g dt$. The normal to $\gamma(t)$, denoted $n(t)$, is defined as the normalized intersection of two planes: perpendicular to $\gamma_t(t)$, and tangent to M on $q = \gamma(t)$.

Let $C = S(\tilde{C}) \subset M$ with $\tilde{C} \subset \Omega$. The perimeter of C is the length of its boundary ∂C . Let γ^C be a parameterized curve mapping $t \in [c, d]$ to the whole of ∂C , then

$$\text{per}(C) = \int_{\partial C} dl = \int_c^d \|\tilde{\gamma}_t^C\|_g dt = L(\gamma^C), \quad (5)$$

where the first equality expresses integration on the boundary regardless of the choice of parameterization, and $\gamma^C = S(\tilde{\gamma}^C)$. The area of C is

$$|C| = \int_C dM = \int_{\tilde{C}} da, \quad (6)$$

where $da = \sqrt{|g|} d\omega_1 d\omega_2$ is often referred to as an *area element*, and the first equality expresses integration on the manifold, regardless of the choice of the parameterization of S with respect to Ω .

Let two functions f_1, f_2 be defined on the manifold M , with their corresponding representations in Ω , u_1, u_2 , respectively, defined in a similar manner to u, f above. Throughout this work we consider functions to lie in a Hilbert space, commonly denoted as $\Gamma_{L_2}(M)$ (see for instance [Güneysu and Pallara 2015]), in which the inner product between two such functions is

$$\int_M f_1 f_2 dM = \int_{\Omega} u_1 u_2 da. \quad (7)$$

Followingly, a gradient operator ∇_g which satisfies $\langle \nabla_g u(\omega_1, \omega_2), \tilde{w} \rangle_g = \lim_{h \rightarrow 0} \frac{f(q+h\tilde{w}) - f(q)}{h}$, $\forall \tilde{w} = J\tilde{w} \in T_q$, with $\|\tilde{w}\|_2 = \|\tilde{w}\|_g = 1$ is obtained. The divergence operator $\nabla_g \cdot$ is then given by the adjoint of ∇_g , i.e. it satisfies,

$$\int_{\Omega} \nabla_g \cdot \tilde{F} u da = \int_{\Omega} \langle \tilde{F}, \nabla_g u \rangle_g da. \quad (8)$$

Finally, we can define the divergence theorem on manifolds: Let $C \subset M$ be a compact set with a smooth boundary ∂C and a boundary normal n^C . Then, C may be treated as a manifold in its own right. Let F be a vector field on C (F is compactly supported), then

$$\int_C \nabla \cdot F dM = \int_{\partial C} F^T n^C dl. \quad (9)$$

We can express (9) also in the parameterization domain⁵:

$$\int_{\tilde{C}} \nabla_g \cdot \tilde{F} d\tilde{a} = \int_{t_1}^{t_2} \langle \tilde{F}, \tilde{n}^C \rangle_g \|\tilde{\gamma}_t^C\|_g dt, \quad (10)$$

where $\tilde{\gamma}_t^C$ is defined as in Eq. (5), as a smooth parametric curve along $\partial \tilde{C}$, i.e. $\gamma^C = S(\tilde{\gamma}^C)$ is along ∂C .

The \mathcal{P} -Laplace-Beltrami is defined as

$$\Delta_{g,\mathcal{P}}(u) := \nabla_g \cdot (|\nabla_g u|^{\mathcal{P}-2} \nabla_g u). \quad (11)$$

For $\mathcal{P} = 2$, $\Delta_{g,\mathcal{P}}(u)$ coincides with the Laplace-Beltrami operator, yielding a diffusion process on surfaces by,

$$\frac{\partial u}{\partial t} = \Delta_{g,2}(u), \quad u(0) = f. \quad (12)$$

Other special cases we will discuss are $\mathcal{P} = 3$, and $\mathcal{P} = 1$.

4 PARAMETRIC SETTING FORMULATION

In this section we present in detail the fundamentals of total variation on parametric surfaces. These serve our theoretical investigations and numerical demonstrations, presented later.

We analyze functions on surfaces in the setting of Sec. 3. Let M be a smooth manifold given as a differentiable and invertible parametric surface $S(\Omega)$, $\Omega \subset \mathbb{R}^2$ with domain variables $\omega_1, \omega_2 \in \Omega$. The Jacobian J maps vectors from Ω to M , inducing the metric g , see Eqs. (2), (4). We assume to have a function $u : \Omega \rightarrow \mathbb{R}$, and a field $z : \Omega \rightarrow \mathbb{R}^2$. We assume that z is differentiable with compact support. Note that u is not necessarily continuous. We examine the following non-Euclidean TV functional,

$$\text{NETV}(u) := \sup_z \int_{\Omega} u \nabla_g \cdot z da \text{ s.t. } \|z\|_g \leq 1 \quad \forall \omega_1, \omega_2 \in \Omega. \quad (13)$$

This is a special case of TV on Riemannian manifolds, reduced to our parametric setting of two-dimensional surfaces (compatible with mesh processing). For the general case, see [Ben-Artzi and LeFloch 2007; Miranda Jr 2003].

$\text{NETV}(u)$ has two notable special cases: The Euclidean metric - which is obtained for a planar M , and the non-Euclidean integral formulation $\int_{\Omega} \|\nabla_g u\|_g da$ - which is obtained for a differentiable function u . In this paper our focus is on non-Euclidean metrics, and non-continuous functions (unless otherwise noted).

Usually, a Neumann boundary condition is assumed, achieving invariance to shift by a scalar,

$$\text{NETV}(u + \alpha) = \text{NETV}(u), \quad \forall \alpha \in \mathbb{R}. \quad (14)$$

Note that if M is closed, then the boundary is empty - hence the Neumann boundary condition holds trivially. Another commonly used property is that NETV is absolutely one-homogeneous, i.e.

$$\text{NETV}(\alpha u) = |\alpha| \text{NETV}(u), \quad \forall \alpha \in \mathbb{R}. \quad (15)$$

⁵This is usually stated in more general setting, for instance see proposition 4.9 in [Gallot et al. 1990]. For usage example see proof of Thm. 1 in [Fumero et al. 2020].

If a field z admits the supremum of $NETV(u)$, Eq. (13), then the field $sign(\alpha)z$ admits $NETV(\alpha u)$. For an in-depth derivation and generalization of basic properties such as the above, see [Miranda Jr 2003], Section 3.

Being absolutely one-homogeneous, we know that the subdifferential of $NETV$, as stated for instance in [Burger et al. 2016], is the following set⁶:

$$\partial NETV(u) = \{p : NETV(v) \geq \int_{\Omega} p v \, da \, \forall v, \, NETV(u) = \int_{\Omega} p u \, da\}, \quad (16)$$

where v is assumed to have the same Neumann conditions as u . By Eqs. (16), (13) we have

$$\nabla_g \cdot z \in \partial NETV(u) \Rightarrow \begin{cases} NETV(u) = \int_{\Omega} u \nabla_g \cdot z \, da. & (17) \\ NETV(v) \geq \int_{\Omega} v \nabla_g \cdot z \, da., & (18) \end{cases}$$

and the converse

$$\int_{\Omega} u \nabla_g \cdot z \, da = NETV(u) \Rightarrow \nabla_g \cdot z \in \partial NETV(u). \quad (19)$$

The $NETV$ minimizing flow, performed on a function $f : \Omega \rightarrow \mathbb{R}$, is defined as

$$u_t = -p(t), \quad p(t) \in \partial NETV(u(t)), \quad u(0) = f, \quad t \geq 0, \quad (20)$$

where $u_t = \frac{\partial u}{\partial t}$. Recently, [Bungert and Burger 2020] proved that eigenfunctions are exposed upon decay by such flows as asymptotic solutions (just before extinction). This enables us to reveal eigenfunctions numerically, by simulating Eq. (20), as done in Figs. 3, 4, 5.

4.1 Indicator Functions

Indicator functions are non-continuous functions which have an important role in total variation analysis. In our non-Euclidean setting we analyze the indicator function of a subset $C \subset M$,

$$\chi^C(q) = \begin{cases} 1 & q \in C \\ 0 & q \in M \setminus C, \end{cases} \quad (21)$$

for which we construct

$$\tilde{\chi}^C = \chi^C \circ S, \quad (22)$$

i.e. $\tilde{\chi}^C$ is the indicator of $\tilde{C} \subset \Omega$ where

$$\tilde{C} = \{S^{-1}(q) \mid q \in C\}. \quad (23)$$

For convenience, let us define the $NETV$ of a set as the $NETV$ of its indicator function $NETV(C) := NETV(\tilde{\chi}^C)$, i.e.

$$NETV(C) := \sup_z \int_{\tilde{C}} \nabla_g \cdot z \, da \, s.t. \, \|z\|_g \leq 1 \, \forall \omega_1, \omega_2 \in \Omega. \quad (24)$$

In the following we assume the setting in which the divergence Thm. on manifolds (10) holds, i.e.: C is a connected set with a smooth boundary ∂C . The boundary has normals n^C (perpendicular to ∂C

⁶The deduced subdifferential needs a product space. We consider the definition of Eq. (7).




Fig. 3. M is a surface-revolution of a part of a translated sinc curve, inducing a non-Euclidean metric. f is initialized as an indicator function of a "sleeve set". Upper row: linear diffusion. Bottom row: $NETV$ minimizing flow. Unlike linear diffusion, the function remains piecewise constant throughout the flow - dividing the surface M to subsets. Iter 50: New boundaries, of small perimeters emerge; iter 100: Initial boundaries subside; iter 150: The sets merge so that only the minimal perimeter remains - which is numerically shown to be an eigenfunction (by Thm. 2.3 in [Bungert and Burger 2020]). When an eigenfunction indicates a set (as is the case here), we call it an eigenset. In the Euclidean case - the eigensets' shape and behaviour have well studied properties [Andreu et al. 2001b; Bellettini et al. 2002]. In contrast, eigensets of the non-Euclidean case are less understood. We introduce novel theoretical properties of the eigensets in non-Euclidean domains.

and tangent to M) with a corresponding \tilde{n}^C s.t. $n^C = J\tilde{n}^C$, where J is the Jacobian from Sec. 3. In addition, we assume a parameterized curve $\gamma^C : t \in [t_1, t_2] \rightarrow \partial C$ for which we construct $\tilde{\gamma}^C = S^{-1} \circ \gamma^C$, a parameterized curve along $\partial \tilde{C}$ ⁷.

Let a field z that is normal to the boundary of C on (almost all) boundary points, and of norm less than or equal to one everywhere on M , i.e.

$$z = \tilde{n}^C \text{ for a.e. } \omega_1, \omega_2 \in \partial \tilde{C}, \quad \|z\|_g \leq 1 \, \forall \omega_1, \omega_2 \in \Omega, \quad (25)$$

where a.e. stands for "almost every". Then z admits the supremum of Eq. (24) for $NETV(C)$. Such a z exists if the boundary of C is differentiable almost everywhere. This condition is satisfied by the assumption of a smooth boundary.

Last property but not least, we have that

$$NETV(C) = per(C). \quad (26)$$

For a proof and further details see Sec. B in the Appendix.

5 THEORETICAL FINDINGS

Here we generalize the theory from the Euclidean setting to closed non-Euclidean manifolds⁸ of the parametric surface setting (Sec. 4). Ultimately a new generalization of convexity is derived, as we prove properties of the eigenfunctions of the sub-gradient. The theory is demonstrated numerically as well.

⁷To understand how S maps boundaries from one domain to another refer to [Lee 2013], Thm. 2.18.

⁸Note that closed manifolds induce a different boundary condition than the non-closed case. The treatment of both cases is similar - for brevity we show the closed case only.

5.1 Derivations

Definition 5.1. u is an eigenfunction of $NETV$ if $\exists \xi : \Omega \rightarrow \mathbb{R}^2$ s.t.

- $\lambda u \in \partial NETV(u)$
- $\nabla_g \cdot \xi = \lambda u$

for some $\lambda \in \mathbb{R}$.

Definition 5.2. C is an eigenset if $\exists \beta > 0$ s.t.

$$\psi := \tilde{\chi}^C - \beta \tilde{\chi}^{M \setminus C} \quad (27)$$

is an eigenfunction of $NETV$, where $M \setminus C$ is the complement of C , and $\tilde{\chi}^{M \setminus C}$ is defined as in Eq. (22).

The existence of eigensets should be verified for each M individually. We provide an example of such verification for a torus in the appendix. We also find eigensets for various manifolds in our numerical experiments.




Fig. 4. $NETV$ minimizing flow on a spherical manifold M . The initial function f assumes two options: An indicator function of a geodesically convex spherical cap C (lower row), or of $M \setminus C$, which is not geodesically-convex (upper row). From right to left: Values of $q_1 \in C$ throughout the flow, Values of $q_2 \in M \setminus C$ throughout the flow, and the flow portrayed as color on M . C and $M \setminus C$ remain intact throughout the flow, and their values change linearly with time t , until they decay completely. Such a behaviour implies that both f s are eigenfunctions, i.e. C and $M \setminus C$ are eigensets. Notably $M \setminus C$ is not geodesically convex. However, in the Euclidean case, eigensets must be convex sets. This raises the question: Is there an alternative notion of convexity on manifolds, other than geodesical convexity, which characterizes eigensets similarly to the Euclidean case? In the following we define and prove such a notion.

CLAIM 1. If C is an eigenset with ψ as in (27) and $\lambda \neq 0$, then

$$\beta = \frac{|C|}{|M \setminus C|}, \quad (28)$$

where $|C|$, $|M \setminus C|$ are areas of C , $M \setminus C$.

PROOF. By manifold divergence theorem (see for instance proposition 4.9 in [Gallot et al. 1990]) we have for closed manifolds that $\int_{\Omega} \nabla_g \cdot z da = 0$, since their boundary is an empty set⁹. ψ is an eigenfunction, i.e. $\exists \xi$ as in Def. 5.1, namely $\nabla_g \cdot \xi = \lambda \psi$. Plugging this we have $\int_{\Omega} \nabla_g \cdot \xi da = \int_{\Omega} \lambda \psi da = 0$. Plugging definition (27) we get

$$\lambda \int_{\Omega} \tilde{\chi}^C - \beta \tilde{\chi}^{M \setminus C} da = 0. \quad (29)$$

⁹See for instance chapter 4.A.2 in [Gallot et al. 2004]

Thus, since $\lambda \neq 0$, by Eq. (6)

$$|C| - \beta |M \setminus C| = 0. \quad (30)$$

□

We note that in case $C = M$, we have by Eq. (29) that $\lambda |M| = 0$, i.e. $\lambda = 0$. This settles well with the following claim.

CLAIM 2. If C is an eigenset, then

$$\lambda |C| = NETV(C). \quad (31)$$

PROOF. Let ψ as in (27). Plugging the eigenfunction property $\nabla_g \cdot \xi = \lambda \psi$ to $\int_{\Omega} \psi \nabla_g \cdot \xi da$ yields

$$\int_{\Omega} \psi \nabla_g \cdot \xi da = \lambda \int_{\Omega} \psi^2 da = \lambda (|C| + \beta^2 |M \setminus C|) = \lambda |C| (1 + \beta), \quad (32)$$

where the last equality uses Eq. (28) as follows: $|C| + \beta^2 |M \setminus C| = |C| (1 + \beta^2 \frac{|M \setminus C|}{|C|}) = |C| (1 + \beta)$. On the other hand, we use Def. 5.1 again, namely $\nabla_g \cdot \xi \in \partial NETV(\psi)$ - which we plug to Eq. (17) and get

$$\int_{\Omega} \psi \nabla_g \cdot \xi da = NETV(\psi) = NETV((1 + \beta) \tilde{\chi}^C - \beta) = (1 + \beta) NETV(C), \quad (33)$$

where the last equality is by Eqs. (15), (14), and the equality before uses $\psi = \tilde{\chi}^C - \beta \tilde{\chi}^{M \setminus C} = (1 + \beta) \tilde{\chi}^C - \beta$. Thus we have by Eqs. (32) and (33) the relation

$$\int_{\Omega} \psi \nabla_g \cdot \xi da = \lambda |C| (1 + \beta) = (1 + \beta) NETV(C), \quad (34)$$

i.e.

$$\lambda |C| = NETV(C). \quad (35)$$

□

COROLLARY. Plugging Eq. (65) to Eq. (31) we have

$$\lambda |C| = NETV(C) = per(C). \quad (36)$$

Definition 5.3. A set $C \subset M$ with a nonempty ∂C is called a *minimal perimeter set* if every set D that contains C has a larger perimeter, i.e.

$$per(C) \leq per(D) \quad \forall D \supset C. \quad (37)$$

Definition 5.4. A set $C \subset M$ with a nonempty ∂C is called a *locally minimal perimeter set* if

$$per(C) \leq \frac{|M \setminus C|}{|M \setminus D|} per(D) \quad \forall D \supset C. \quad (38)$$

This definition is a weaker than the previous definition, since $\frac{|M \setminus C|}{|M \setminus D|} > 1$. The locality of this definition can be explained as follows: The larger $\frac{|M \setminus C|}{|M \setminus D|}$ is, the further away some points in D must be from the local neighbourhood of C . Noteable cases:

- Any minimal perimeter set is also a locally minimal perimeter set.
- Closed surfaces have $per(M) = 0$ thus they do not contain minimal perimeter sets, but they do contain **locally** minimal perimeter sets.
- In the Euclidean case, both these definitions coincide with the set being convex.

THEOREM 5.5. Any eigenset C , a subset of a closed M , satisfies

$$NETV(C) \leq \frac{|M \setminus C|}{|M \setminus D|} NETV(D), \quad \forall D \supset C. \quad (39)$$

PROOF. By the eigenset Def. 5.1, $\exists \psi$ as in (27) with a field ξ s.t. $\nabla_g \cdot \xi \in \partial NETV(\psi)$. Thus by Eq. (18) we have

$$\int_{\Omega} \tilde{\chi}^D \nabla_g \cdot \xi \, da \leq NETV(\tilde{\chi}^D), \quad (40)$$

where $\tilde{\chi}^D$ is defined similarly to $\tilde{\chi}^C$. By definition of $\tilde{\chi}^D$ we have $\int_{\Omega} \tilde{\chi}^D(\cdot) \, da = \int_{\tilde{D}}(\cdot) \, da$. Using this and the notation of Eq. (24), we have

$$\int_{\tilde{D}} \nabla_g \cdot \xi \, da \leq NETV(D). \quad (41)$$

On the other hand, we have the eigenset property $\nabla_g \cdot \xi = \lambda \psi$. Plugging this to $\int_{\tilde{D}} \nabla_g \cdot \xi \, da$, we have

$$\int_{\tilde{D}} \nabla_g \cdot \xi \, da = \lambda \int_{\tilde{D}} \psi \, da = \lambda(|C| - \beta(|D| - |C|)). \quad (42)$$

Note that we integrate over D , while ψ is defined for C, M . The last equality uses $C \subset D$ which translates to $\tilde{C} \subset \tilde{D}$ by invertability of S . Thus by Eqs. (41), (42)

$$\lambda(|C| - \beta(|D| - |C|)) \leq NETV(D). \quad (43)$$

Plugging Eqs. (31), (28) we get

$$\frac{NETV(C)}{|C|} \left(|C| - \frac{|C|}{|M \setminus C|} (|D| - |C|) \right) \leq NETV(D), \quad (44)$$

$$NETV(C) \frac{|C|}{|C|} \left(\frac{|M \setminus C| - (|D| - |C|)}{|M \setminus C|} \right) \leq NETV(D), \quad (45)$$

using $|M \setminus C| - (|D| - |C|) = |M| - |C| - |D| + |C| = |M \setminus D|$, we finally have

$$NETV(C) \leq \frac{|M \setminus C|}{|M \setminus D|} NETV(D). \quad (46)$$

□

COROLLARY. Plugging Eq. (65) to Eq. (46) we obtain

$$per(C) = NETV(C) \leq \frac{|M \setminus C|}{|M \setminus D|} NETV(D) = \frac{|M \setminus C|}{|M \setminus D|} per(D) \quad \forall D \supset C, \quad (47)$$

thus any eigenset C , subset of a closed M , is a locally minimal perimeter set.

If $|M| = \infty$, then C is a minimal perimeter set. This happens since locally minimal perimeter subsets of surfaces with $|M| = \infty$ automatically become minimal perimeter, by $\frac{|M \setminus C|}{|M \setminus D|} = 1$. Furthermore, if M is Euclidean (with $|M| = \infty$), then this corollary shows that C is convex - which is a well known property for Euclidean TV . Hence, the new notion of locally minimal perimeter is a generalization of set convexity, with regard to total-variation theory.




Fig. 5. Weighted indicator functions on Torus and their $NETV$ flow. First row: Two geodesical disks. The small disk is defined in "Euclidean position", i.e. results for Euclidean case carry over. Namely - the small disk is both geodesically convex, and a (locally) minimal perimeter set. The large disk is not in Euclidean position, and Euclidean laws do not carry over: It is not geodesically convex, nor is it of locally minimal perimeter. Throughout the flow it becomes a sleeve-like set, which is a minimal perimeter set, after which it decays completely. Bottom row: Indicator function of two sleeve sets C_1, C_2 , weighted by the ratio of their areas: $f = \tilde{\chi}^{C_1} + \frac{|C_2|}{|C_1|} \tilde{\chi}^{C_2}$. The sleeve sets are eigensets, as shown in the appendix. By Eq. (36), and since $per(C_1) = per(C_2)$, the eigenvalue ratio should be inverse the area ratio i.e. $\frac{\lambda_2}{\lambda_1} = \frac{|C_1|}{|C_2|}$. Thus we can reformulate as $f = \tilde{\chi}^{C_1} + \frac{\lambda_1}{\lambda_2} \tilde{\chi}^{C_2}$. If this is true, then by Eq. (51) we expect the sets to completely decay at the same time - and indeed they do.

5.2 Overview: Application of the Theory

Eigenfunctions of the non-Euclidean total variation functional are stable, linearly decaying, modes of the minimizing flow. They are expected to be seen throughout the flow, especially upon convergence [Gilboa 2013], [Bungert and Burger 2020]. Fig. 4 demonstrates such linear decay - indicating that the decaying function is an eigenfunction. Such eigenfunctions divide the non-Euclidean manifold to subsets, called "eigensets". For the Euclidean setting it is well established that such eigensets must be convex - giving interpretability and understanding of the expected behaviour of the total-variation minimizing flow, and its regularization properties in general [Andreu et al. 2001b], [Bellettini et al. 2002]. Until now, a fitting generalization was not proven for the non-Euclidean setting.

It was recently proposed [Fumero et al. 2020], based on empirical evidence, that eigensets of the non-Euclidean total variation are geodesically convex. However, Fig. 4 demonstrates numerically a linearly decaying set, i.e. an eigenset, that is not geodesically convex.

In our theoretical derivations we begin by finding the expected values of the eigensets. These enable us to find the eigenvalue associated to the eigenfunction. The eigenvalue readily gives us understanding of the flow behaviour - as it is equal to the rate of linear decay of the eigenfunction. Fig. 5 demonstrates numerically that the rate of decay is as expected. Finally - we derive a necessary condition for a set to be an eigenset, which we call the "locally minimal perimeter" condition. This condition is validated to be a non-Euclidean generalization of the well known Euclidean set convexity property. In Fig. 3, we observe the locally minimal perimeter

criterion on the stable modes of the flow. The influence of the locally minimal perimeter criterion can also be seen across our shape processing experiments which follow, most notably in Figs. 6, 12, 13, 14, 20.

6 SPECTRAL DECOMPOSITIONS

Here we show straightforward extensions of some of the observations done in [Burger et al. 2016] to our settings. Let X be a space of functions on a surface M equipped with a metric g , as described in Eqs. (2), (4). Let $p : X \rightarrow X$ be a zero-homogeneous operator, i.e.

$$p(\alpha f) = \text{sign}(\alpha)p(f), \quad \alpha \in \mathbb{R}, f \in X, \quad (48)$$

with $p(0) = 0$. We examine the following flow:

$$u_t = -p(u(t)), \quad u(0) = f \in X, \quad t \geq 0, \quad (49)$$

where $u_t = \frac{\partial u}{\partial t}$. We assume the flow exists and that the solution is unique. Unlike [Burger et al. 2016], here p maps functions on non-Euclidean domains. Nevertheless, the time domain, denoted by t , is Euclidean. Note that elements in $\partial NETV$ are zero-homogeneous, thus, $NETV$ flow (subgradient descent of the energy) is a zero-homogeneous flow. We also assume the second time-derivative of u exists in the distributional sense almost everywhere and define $\phi : t \rightarrow X$ as

$$\phi(t) = t \cdot u_{tt}. \quad (50)$$

Let ψ be an eigenfunction with respect to p with a positive eigenvalue, i.e. $\exists \lambda \in (0, \Lambda < \infty) : p(\psi) = \lambda\psi$. Let $f = \psi$, then the solution of Eq. (49) is

$$u(0) = \psi \Rightarrow u(t) = \begin{cases} (1 - \lambda t)\psi & t \leq \frac{1}{\lambda} \\ 0 & t > \frac{1}{\lambda} \end{cases}. \quad (51)$$

This can be verified by having, for $t < \frac{1}{\lambda}$, the relation

$$p(u(t)) = p((1 - \lambda t)\psi) = p(\psi) = \lambda\psi = -u_t, \quad (52)$$

where the second equality uses 0-homogeneity of p , the third equality uses the eigenfunction property, and the last equality is an evaluation of u_t by (51). For $t = \frac{1}{\lambda}$ we have a steady state since $p(0) = 0$. By uniqueness of the solution we are done.

We note that since we are examining smoothing processes, p in general is a positive semidefinite operator, $\langle f, p(f) \rangle \geq 0, \forall f \in X$. Thus the eigenvalues are positive. In the case of negative eigenvalues, the flow diverges (but for a finite stopping time can still have a solution).

Thus, for eigenfunctions of positive eigenvalues we get $\phi(t) = \delta(t - \frac{1}{\lambda})f$, i.e. ϕ 's energy is concentrated in a single scale ("frequency") which corresponds to the eigenvalue of f , $\lambda = \frac{1}{t}$. For a general $f \in X$, this motivates the interpretation of ϕ as a spectral transform of f , where the spectral components are positive eigenfunctions of p , in a similar manner to [Bungert et al. 2019; Burger et al. 2016; Gilboa 2014b].

We can compute the reconstruction formula, for a general stopping time T , using integration by parts (and assuming $u_t(0)$ is bounded), by $\int_0^T \phi(t) dt = tu_t|_0^T - \int_0^T u_t dt = Tu_t(T) - u(T) + f = -Tp(u(T)) - u(T) + f$. Denoting the residual $R = Tp(u(T)) + u(T)$, the following reconstruction identity holds $f = \int_0^T \phi(t) dt + R$. This

gives rise to a filtered reconstruction via a filter $H : t \rightarrow \mathbb{R}$ as follows:

$$f^{\text{filtered}} = \int_0^T H(t)\phi(t) dt + R. \quad (53)$$

An important special case is the operator $p \in \partial NETV$, which is zero-homogeneous. Hence a $NETV$ -minimizing flow can be used to filter signals by (53). Moreover, in the case of an initialization with a single eigenfunction we expect a linear decay. In Figs. 5, 6 we use these properties to demonstrate some of our theoretical findings from Sec. 5. In the following we demonstrate how this can be carried over to shape processing.

7 SHAPE PROCESSING

We suggest three methods for nonlinear filtering of shapes, in the framework of Sec. 6. The methods differ by the choice of the operator p of the respective flow. Each method is inspired by a different flow: $M1$ by the Heat Flow, $M2$ by cMCF and $M3$ by MCF. $M3$ uses $p \in \partial NETV$, hence theory of Sec. 5 applies. $M1$ and $M2$ use different yet related operators, and theory regarding these is left as future work. Nevertheless - in all three methods we attain good feature control via manipulation of the spectral components. We will also demonstrate how the choice of different operators p induces different qualities.

So far, we processed a function f on a 2D manifold M embedded in 3D, via Eqs. (49), (50), (51), (53). Here we wish to process the manifold itself - and for this purpose we will choose a function f that describes M . In the setting of parameterized surfaces, the surface function may be the function of choice, i.e. $f = S$, a vectorial function with three channels as the three coordinate functions $x_0(\omega_1, \omega_2), y_0(\omega_1, \omega_2), z_0(\omega_1, \omega_2)$. Note that S also induces the intrinsic metric g . This choice is widely used for shape flows, thus it enables a comparison between our framework and the classical ones. Other representations can be used, see for example Fig. 7. Denote the evolving shape at time t as $S(t)$, and denote $c(t)$ as any evolving coordinate function of $S(t)$, that is $c(t)$ may assume $x(t), y(t)$, or $z(t)$.

7.1 Modifying flows for nonlinear spectral processing

Our framework requires a zero-homogeneous flow evolving on a fixed metric, which is required to induce spectral linear decay of the eigenfunctions. Denote the metric induced by the initial shape g_0 . It is fixed throughout the flow. Denote the evolving shape's metric g_t , which changes throughout the flow, i.e. it is not fixed.

Examining Heat Flow, MCF and cMCF we find that none of these flows is zero-homogeneous, and Heat Flow is the only one performed on a fixed metric. Hence, adaptations of these flows are required.

7.2 Naive method: Unpaired Coordinate Spectral TV

The naive approach utilizes a modification of Heat Flow for our framework. Heat Flow processes each coordinate function independently via Eq. (12), utilizing the Laplace-Beltrami on the fixed metric g_0 throughout the flow. Thus it satisfies a fixed metric, but it is not zero-homogeneous, and a modification is required. By replacing the Laplace-Beltrami with the 1-Laplace-Beltrami of Eq.




Fig. 6. Upper row: M is the surface seen also in Fig. 3. f is initialized as the z coordinate function, shown in color on the shape. A zero-homogeneous operator, $\partial NETV$, induces the spectral components - used for filtered reconstruction. Middle row depicts filtered coordinate function as color on shape. Bottom row: Having a geometrical meaning - the filtered coordinate function can replace the shape's coordinate - Bottom row transforms middle row in this manner, followed by a simple graph cut for clearer visibility of vertex location. The LPF (low pass filter) H_1 exposes an underlying geometrical structure of three planes - corresponding to three "sleeve sets". The sleeve sets decay linearly throughout the flow, proving to be spectral components of $\partial NETV$. The geometrical meaning of linear decay is constant velocity. H_2 further filters and shrinks the upper two planes to one - showing that the bottom plane is of lowest eigenvalue, as it is the slowest decaying plane. H_0 shrinks the three underlying planes to a single one, the residual plane. H_3 adds the non-planar structures to the residual plane. The naive choice of coordinate function is not always natural to the processed shape.

(11) zero-homogeneity is achieved, which results in the operator $-p_{Naive}(c) := \Delta_{g_0,1}c$, and a per-coordinate flow is defined by setting $p(u(t)) = -p_{Naive}(c(t))$ in Eq. (49). Each channel evolves separately, hence the name "unpaired coordinates". We can now perform nonlinear spectral filtering as in Eq. (53), demonstrated on the meteor model in Fig. 8. Note the axis squaring effect, which violates rotation invariance, and also restricts the underlying structure, resulting with bad separation between structure and detail. With that said, it does provide relatively aesthetic results, if one desires such squaring effect.

7.3 Method 1 (M1): Shape Spectral TV

Here we take into account shape coordinate inter-correlations, i.e. we go from coordinate to shape processing. We apply a vectorial flow (as in VTV) on meshes, which results in the operator,

$$-p_{M1}(c) := \nabla_{g_0} \cdot \left(\frac{\nabla_{g_0} c}{\sqrt{\sum_{\tilde{c}=x,y,z} |\nabla_{g_0} \tilde{c}|^2}} \right). \quad (54)$$

Note that the metric is fixed as g_0 . We can also verify that the operator is zero-homogeneous.

Remark: Similar flows were proposed in the past. One important example is [Elmoataz et al. 2008], from which we borrow the combined gradient magnitude of the denominator - designed to account for the inter-correlation of the coordinate function. With that said, the gradient and divergence operators used in [Elmoataz et al. 2008] are significantly different - as they are obtained for general graph structures, while we use operators which account for the surface properties of the mesh. For instance, the graph gradient used in [Elmoataz et al. 2008] is calculated per-vertex and has a non-fixed dimensionality that is equal to the number of edges connected to said vertex, while the surface mesh gradient we use is obtained per-triangle, and is parallel to said triangle, i.e. it is embedded in \mathbb{R}^3 .

The flow is followed by per-coordinate spectral processing as in Eq. (50), (53) - thus x, y, z inter-correlation is preserved. Compared to the naive approach, $M1$ preserves better the initial underlying structure, as demonstrated in Fig. 8. Good multi-scale feature control is demonstrated as well in Fig. 9.

7.4 Method 2 (M2): Conformalized 3-Laplace

Here we modify cMCF [Kazhdan et al. 2012] to our framework, by presenting a conformalized \mathcal{P} -Laplace, as described below. Our




Fig. 7. Handling real world shapes, demonstrated on scan of the famous archeological statue "Bust of Queen Nefertiti". In this example, geometric details are captured as a function that represents perturbation from an underlying smooth shape (further details in Sec. 7.5). In color we see these geometric details experiencing a *NETV* minimizing flow. To obtain the result of Fig. 2, we follow a procedure similar to the one presented in Fig. 6: Spectral representations are obtained from the flow and used for filtering. Followingly the original geometrical details are replaced by their filtered versions, which results with the final filtered shapes.

flow inherits cMCF's limb-head smoothing capabilities (Fig. 10), which we then use for shape filtering. The metric of cMCF is $\tilde{g}_t = \sqrt{|g_0^{-1}g_t|g_0}$, is not fixed. The operator driving the flow, the conformalized Laplace-Beltrami, $\sqrt{\frac{|g_0|}{|g_t|}}\nabla_{g_0} \cdot \nabla_{g_0}$, depends on the evolving shape's metric g_t . To achieve a fixed metric, we re-interpret $|g_t|$ as an operator on the fixed metric g_0 . This is valid since the diffused shape defines both the diffused function as well as the evolving metric. This affects homogeneity, as shown below. We define the conformalized \mathcal{P} -Laplace as,

$$\tilde{\Delta}_{g,\mathcal{P}}(c) := \sqrt{\frac{|g_0|}{|g_t|}}\nabla_{g_0} \cdot (|\nabla_{g_0}c|^{\mathcal{P}-2}\nabla_{g_0}c). \quad (55)$$

By Eq. (4) we have that $|g_t|$ is absolutely 4-homogeneous, hence $\tilde{\Delta}_{g,\mathcal{P}}$ is $\mathcal{P}-3$ homogeneous, $\tilde{\Delta}_{g,\mathcal{P}}(\alpha c) = \sqrt{\frac{|g_0|}{|\alpha|^4|g_t|}}\nabla_{g_0} \cdot (|\alpha|^{\mathcal{P}-2}|\nabla_{g_0}c|^{\mathcal{P}-2}\alpha\nabla_{g_0}c) = \frac{\alpha}{|\alpha|^{\mathcal{P}-3}}\tilde{\Delta}_{g,\mathcal{P}}(c)$. Thus we choose $\tilde{\Delta}_{g,3}$ as a zero-homogeneous modification of the conformalized Laplace. Once again inter-correlations are accounted for, as in [Elmoataz et al. 2008], yielding the operator,

$$-p_{M2}(c) := \sqrt{\frac{|g_0|}{|g_t|}}\nabla_{g_0} \cdot \left(\sqrt{\sum_{\tilde{c}=x,y,z} |\nabla_{g_0}\tilde{c}|^2} \nabla_{g_0}c \right) \quad (56)$$

The flow is followed by nonlinear spectral filtering, Eq. (53). Editing extremities, a capability inherited from our conformal 3-Laplace flow, is demonstrated in Fig. 11, where extremities are in the form of human limbs and head.

7.5 Method 3 (M3): Directional Shape TV

Mesh TV smoothing typically preserve pointy surface points, e.g. tip of chin [Fumero et al. 2020] or ears [Elmoataz et al. 2008]. Here




Fig. 8. A real-world 3D scan of a Meteorite¹⁰. In this example, the shape is described by its coordinate functions. Unlike linear shape processing, the nonlinear case requires special attention to deal with the correlation between the X , Y and Z coordinate functions. In Method 1 (M1) such considerations are accounted for. From top left, counter clockwise: 1) original model; 2) M1 low-pass filter (LPF); 3) Naive method low-pass filtering, ignoring the inter-coordinate correlation, this method has an axis-squaring effect. Even such harsh filtering with the naive method dose not yet fully smooth out details, suggesting lack of feature control. Contrary, M1 smooths out details without distorting underlying structure, as expected from a smoothing procedure; 4) Details obtained via M1 high-pass filter (HPF) are added to the model after squaring it with the naive method LPF

we propose a method that preserves edges, e.g. muscle contour, similarly to TV processing of images. While M1 and M2 utilized modifications of Heat Flow and cMCF, M3 draws inspiration from MCF.

MCF already has a thoroughly researched fixed-metric zero homogeneous modification: The TV flow as applied to gray-scale images [Kimmel et al. 2000]. For a surface represented as $S = (x, y, f(x, y))$, this modification entails constraining the evolved shape to be of the form $S(t) = (x, y, f(x, y, t))$. This is enforced by constraining each point on the surface to evolve in direction \hat{z} (perpendicular to the x, y plane). We note that unconstrained MCF would necessarily violate this form of $S(t)$, as it theoretically converges to a singular point.

¹⁰Edmore Meteorite 3D scan, courtesy of Prof. Jeremy Davis, Central Michigan University, Department of Art and Design

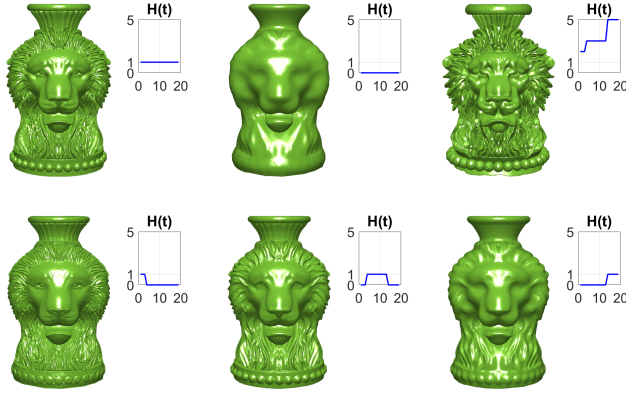


Fig. 9. M1 filtering. Top row, left to right: All-pass reconstruction; residual; 3-bands amplification. Bottom row: each of the 3 bands isolated on top of the residual. As t grows, choosing $H(t)$ greater or less than 1, results in amplification or attenuation of coarser details, as can be observed e.g. in the model's contour and clumping of hair-stands.




Fig. 10. Three upper rows: M2 conformalized \mathcal{P} -Laplace flows. $\mathcal{P} = 2$ is an unscaled version of the conformalized Mean Curvature Flow (cMCF). $\mathcal{P} = 1$ is a new conformalized shape TV flow. For $\mathcal{P} = 3$ the flow is zero-homogeneous. Bottom row: cMCF from [Kazhdan et al. 2012]




Fig. 11. Shape exaggeration by M2 spectral filtering. The filter inherits its properties from the conformalized 3-Laplace flow (Fig. 10), which translates to interesting limb-head editing. Isometry robustness is demonstrated as well.

Our third method aims to generalize the above direction-constraint to general shapes, hence the name "directional". The x, y domain is generalized to be an over-smoothed version of the initial shape which we denote \hat{S} . Each $p \in S$ is mapped to a $\hat{p} \in \hat{S}$. The direction of evolution is fixed as $\hat{d} = \alpha \frac{S - \hat{S}}{|S - \hat{S}|}$, where α is a sign indicator which ensures \hat{d} points "outwards". Finally, the evolving initial surface is represented as $f = \alpha|S - \hat{S}|$. Note that $S = \hat{S} + f\hat{d}$. This method is a generalization in the following sense: Consider the form $S = (x, y, f(x, y))$, choosing $\hat{S} = (x, y, 0)$, we have that $\hat{d} = \hat{z}$, and $\alpha|S - \hat{S}| = (0, 0, f(x, y))$.

The notion processing d This method belongs to a vast family of shape processing techniques, operating on the displacement field

We advocate the choice of \hat{S} as a cMCF smoothed version of S , since cMCF was shown to provide a conformal mapping from S to \hat{S} . By construction - the metric is fixed and inter-correlations are accounted for. The proposed zero-homogeneous operator (acting on a scalar-valued function u) is,

$$-p_{M3}(u) := \nabla_{g_0} \cdot \frac{\nabla_{g_0} u}{|\nabla_{g_0} u|}, \quad (57)$$

where $u(t)$ is the evolution of f at time t , which results in $\frac{\partial S}{\partial t} = \nabla_{g_0} \cdot \frac{\nabla_{g_0} u}{|\nabla_{g_0} u|} \hat{d}$, satisfying the imposed directionality. Finally, f is filtered as in Eq. (53), and a filtered shape is obtained by $\hat{S} + f^{\text{filtered}} \hat{d}$. Being closely related to spectral TV on images, this method preserves detail well, as demonstrated in Fig. 12. Though inspired by MCF, the proposed flow is substantially different.

8 SHAPE DEFORMATION

Here we show a novel application. For the first time we perform total-variation for the shape deformation task. We will see that this induces a piecewise-constant deformation field, where the deformation concentrates on small-perimeter boundaries, showcasing the eigenset properties derived in Sec. 5.

The shape deformation problem involves finding a plausible transformation of a given shape while accommodating constraints. The constraints are specified by the user as points, or regions of the shape, which are to be moved away from their original location. The deformation process should ensure that the resulting shape maintains its structural integrity while sufficing the constraints.

To this end we propose a constrained minimization of the total-variation of the displacement field, defined as

$$d' := S' - S, \quad (58)$$

where S is the coordinate function of the original given shape, and S' is the resulting shape. The minimization process is provided below as Algorithm 1. The algorithm uses a minimization process proposed by [Bronstein et al. 2016] for minimizing L1 norms on manifolds. A slight difference is, that we use a vectorial version of [Bronstein et al. 2016]. The minimizing operator of choice is p_{M1} of Eq. (54). Minimization is performed while enforcing the deformation constraints. As an initial solution to the process, we use the gradient-based linear deformation proposed in [Botsch and Sorkine 2007]. We use a matrix version of p_{M1} , denoted p_{M1}^{matrix} , which is updated at each iteration. Further details are available in Appendix D.




Fig. 12. Upper row: Low-pass filtering applied for smoothing. All our methods can utilize rough time discretization for runtime efficiency at the expense of reconstruction error. In this example, M3 is used, and completed in approximately 30 seconds. In contrast, the Laplace-Beltrami method involves SVD for 2000 eigenvectors of a $[17 \cdot 10^3 \times 17 \cdot 10^3]$ sparse matrix, requiring about 5 hours (2.5 orders of magnitude slower). Nonetheless, more efficient linear methods exist, such as in [Cignoni et al. 2005], which take the same amount of time. Bottom row: Smooth caricaturization via bandpass exaggeration coupled with smoothing. We compare the exaggeration capabilities of M3 with those in [Cignoni et al. 2005], which also exaggerates shapes using the deformation from an over-smoothed underlying shape. Bottom left (zoom-in): Our approach effectively smooths smaller details, such as the scales of the armadillo, while maintaining larger structures like the knees. Our theoretical findings correlate "resistance to smoothing" of geometrical structures with generalized convexity and the ratio of perimeter to area. The knees, being more convex-like with a lower perimeter to area ratio than the scales, indeed demonstrate greater resistance to smoothing. This results in superior separation of detail compared to both Laplace-Beltrami smoothing and the shape exaggeration method in [Cignoni et al. 2005]. For additional reference, exaggeration without smoothing is presented in the appendix, Fig. 21, where we also demonstrate consistency with the caricaturization principles discussed in [Sela et al. 2015].

As illustrated in Fig. 13, the displacement field d' gradually becomes piecewise constant during the shape deformation process. Interestingly, the deformation tends to concentrate on boundaries with small perimeters, suggesting a potential connection to eigenfunctions. We validate this hypothesis numerically in Fig. 14 by analyzing $\|d'\|_g$, the magnitude of the field. See Fig. 15 for method comparison. For completeness, we added additional results in the appendix, Fig. 19.

9 DISCUSSION

9.1 Non-differentiable shapes: A limitation of our theory

Our theoretical findings rely on the notion of "good" metric spaces [Miranda Jr 2003]. In the context of our shape processing applications, these metrics are defined by the specific shapes being processed. However, within the domain of computer graphics, the

Algorithm 1 Shape Deformation

```

 $S \leftarrow$  original unconstrained shape coordinate function
 $b \leftarrow$  boundary indices
 $\kappa \leftarrow$  boundary constraints
 $w \leftarrow \infty$  ▷ A high constraint weight
 $S' \leftarrow$  initial solution
while True do
   $S'^{cached} \leftarrow S'$ 
   $\partial VNETV \leftarrow$  construct the minimizing operator considering  $S, S'$ 
   $S' \leftarrow \text{minimize}(\|\partial VNETV(S' - S)\|_2^2 + w \cdot \|S'(b) - \kappa\|_2^2)$ 
  if  $S'^{cached} \approx S'$  then ▷ Convergence stopping condition
    break
  end if
end while
return  $S'$ 

```




Fig. 13. Our TV regularization process of the displacement field in the shape deformation setting for Snail, Armadillo, Cheburashka and Knubble models. The deformation constraints, which are all visible here, are marked with red dots. All cases show the tendency for a piecewise constant deformation field. The sets of constant deformation inherit their properties from the eigensets of Sec. 5, namely - their boundaries have small perimeters, e.g., neck and elbows. Note in the zoomed-in windows (Second row, most-right), how the not-stretched parts remain intact.

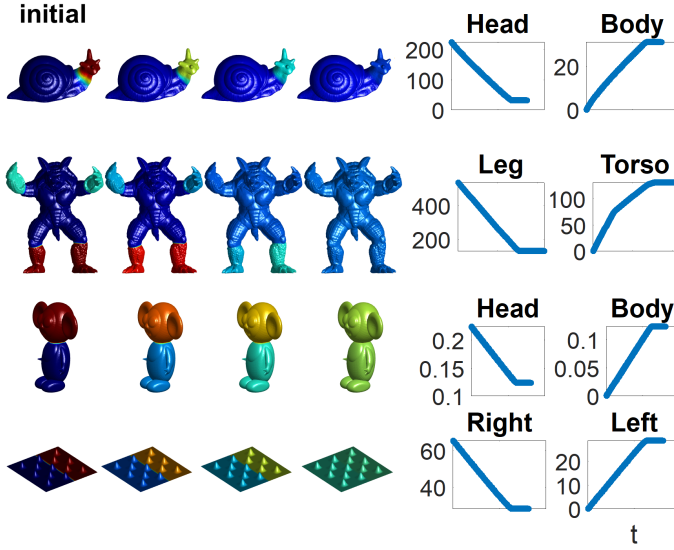


Fig. 14. The magnitude of our resulting deformation field d' tends to be constant on *NETV* eigensets, as validated here numerically. Similarly to Figs. 4, 5, we check whether a set is an eigenset by testing whether it is a linearly decaying mode of the *NETV* flow. Left Column: After obtaining the deformation (right column in Fig. 13), we calculate the magnitude of the deformation field vectors, $\|d'\|_g$. We visualize these magnitudes using color, and it can be observed that they exhibit a piecewise-constant pattern. Middle Columns: An *NETV* minimizing flow is applied to the magnitudes $\|d'\|_g$. The sets of constant value remain intact until they completely decay. Right Columns: Here the numerical values of the sets are displayed throughout the flow. A linear decay can be observed - implying that these are indeed eigensets. Observations: 1) In the snail case, the initial $\|d'\|_g$ slightly deviates from the typical crisp boundaries observed in other shapes. Nonetheless, clearer boundaries appear early-on in the flow, and the linear decay quickly becomes similar to the other shapes. 2) The decay of the Armadillo's torso is different since it exhibits two distinct rates. This can be explained by our theoretical findings: One rate happens before the arm sets coincide with the torso set, and another after. Following the coincidence of the arms and torso sets, the area increases while the perimeter decreases due to the disappearance of the arm perimeters. This leads to a reduction in the eigenvalue λ according to Eq. (36), resulting in a decrease in the decay slope - as evident in the observed behavior and as anticipated by Eq. (52).

shapes encountered can often exhibit high non-differentiability, which may even be enhanced by discretization. Consequently, our assumptions may not be suitable in such cases. While works concentrating on such considerations where conducted for the Laplacian operator (see for instance [Wardetzky et al. 2007], [Sharp and Crane 2020]), this remains for future work regarding the operators used in this paper.




Fig. 15. A shape-deformation method comparison for the Snail¹¹ model, considering the same constraints in all methods. Upon using our method for shape deformation with total-variation, the stretch is concentrated on boundaries of plausible shape segments. In contrast - linear methods stretch the shape globally, and As-Rigid-As-Possible methods (ARAP) globally penalize the local stretch. Other comparisons are available at Fig. 19.




Fig. 16. Time step size of the discrete flow is an important parameter, controlling a trade off between flow accuracy, and execution complexity. An all-pass filter should perform reconstruction (in all our methods), however too large time steps result with inaccurate flow, which yields inaccurate spectral representations and reconstruction. Here we see all-pass filters of the Armadillo model using increasingly bigger time steps ($M1$). Reconstruction error is shown in color per-vertex, and is larger with time step size, as expected.




Fig. 17. Image Spectral TV (upper row) and Spectral ATV (bottom row) filtering, implemented as special cases of $M3$. Unlike the usual Spectral ATV, here the metric is induced by an explicit surface given by the image function (bottom left). Left-most column: Low-pass filter (LPF), demonstrating the difference between the "convexifying" effect of Euclidean spectral TV compared to the minimal perimeter effect of the non-Euclidean spectral *NETV* on surfaces. Middle column: Band stop filter (BSF), namely high-pass details where added to the LPF, resulting with a natural looking image

¹¹Snail model, courtesy of 3Demon team. Licensed under the Creative Commons - Attribution - Share Alike.




Fig. 18. Representations that enable good feature separating filters are informative, hence may play a role in other tasks. Here we demonstrate an initial attempt at using $M2$'s spectral representation for co-segmentation of the Michael models in various poses: Each point (vertex) on each pose (mesh) is assigned the representation $\phi(t)|_{q \in M}$. For co-segmentation we partitioned representation space using all representations from all poses - yielding segments that are shared across poses. Space partitioning was done via the k-means algorithm, using earth-movers distance instead of Euclidean distance. Earth-movers is a more appropriate measure since it is better at quantifying perturbations along the t dimension.

9.2 Flow-induced spectral representation considerations

Our filtering framework is flow-based, where spectral representations are manifested as linear decaying components. Shape analysis is performed along the time domain. In contrast, the Laplace-Beltrami eigenfunctions are acquired by solving an eigenvalue problem on the non-Euclidean domain. One advantage of the flow-based framework is computational: The numerical simulation of a shape flow is often computationally cheaper than a numerical solution to an eigenvalue problem performed on the shape domain (e.g. solving the SVD of a discrete Laplace-Beltrami operator) - as demonstrated in Fig. 12. This results with better filtering complexity. This is opposed to the Euclidean case - where eigenvalue decomposition of the Laplacian is not needed, since it is fixed and known (Fourier basis). Simulating a flow requires discrete time steps, resulting with a trade-off between computational complexity and simulation accuracy, controlled by the time step size. This is true for Euclidean and non-Euclidean settings alike. To measure the accuracy of the flow, we can test the reconstruction error of an all pass filter, i.e. $H(t) = 1 \forall t$, as demonstrated in Fig. 16. Remark: Our semi-implicit implementation of the flows exhibits stability for large time-steps, even when they result with inaccuracies. This is attributed to the time-steps of both the filtering flows and the deformation flows being approximated as solutions to stable L2 minimization problems (similar stability may be found for instance in [Kazhdan et al. 2012]).

Another consideration arises when comparing two representations, $\phi_1(t), \phi_2(t)$. The usual L_2 distance may not be compatible on

time domains, and even more-so on discretized time domains. This is since L_2 does not express well the difference between small and large shifts in time. For example, consider an eigenfunction - which was shown to have a representation $\phi(t) = \delta(t - \frac{1}{\lambda})$, where λ is the eigenvalue. Thus the L_2 measure will not be able to discriminate between eigenfunctions with similar eigenvalues to eigenfunctions with largely different eigenvalues. Thus appropriate distance measure must differentiate small from large time perturbations. One such measure is the earth movers' distance. An example using this distance on ϕ is portrayed in Fig. 18, where co-segmentation takes place.

9.3 Spectral image TV and shape spectral TV relation

Our shape spectral TV methods require a zero-homogeneous flow performed on a fixed metric. Considering the Beltrami and TV flow equivalence presented in [Kimmel et al. 1998] (see a brief reminder of this equivalence in the appendix, Sec. E), we have that applying spectral TV to images is a form of shape spectral TV. Let us have a closer look at this statement: To transition from MCF to the equivalent TV flow, the flow was rephrased on a fixed metric (the Euclidean pixel grid), which was absorbed as a nonlinearity of the operator, making it zero-homogeneous. Thus all requirements of the zero-homogeneous spectral framework were met.

With that said, this framework is more restrictive than our general framework in 3 ways: The operator suggested is one of a kind, the

shape has to be parameterized as an image function, and the fixed metric is a Euclidean domain (the pixel-grid).

In *ATV* the latter constraint is not required. As shown in [Biton and Gilboa 2022], *ATV* can be re-interpreted using a generalization of the gradient, $\nabla_A f = A(x)\nabla f$. This gradient generalization can be obtained by considering the pixel grid as the Ω domain in Eq. (2), and A as the metric g from Eq. (4), induced by some unspecified M . While other aspects of *ATV* do not coincide with the differential geometry framework we use here, it is certainly related to our work. Interestingly, in *ATV* the importance of parameterization domain is greater than in our framework, as the signal lies in Ω , and gradients are on Ω as well, mapped from an unspecified non Euclidean domain. In contrast, our signal lies on an explicit manifold M .

M3 generalizes both spectral *TV* and *ATV* in the following sense: Considering the form $S = (x, y, f(x, y))$, choosing $\hat{S} = (x, y, 0)$, we have that $\hat{d} = \hat{z}$, and $\alpha|S - \hat{S}| = (0, 0, f(x, y))$. Now consider two options - option 1: \hat{S} induces g , resulting with a Euclidean flow of an image function. Option 2: S induces g , resulting with a non Euclidean flow of an image function on an adapted metric. This is a form of *ATV*, but with a specified non-Euclidean surface domain. See example in Fig. 17.

9.4 Future Ideas

Our flow-based spectral framework can easily be adapted to a wide collection of operators that assume the required homogeneity and a fixed metric. Inevitably - neural-networks come to mind, where homogeneity can be taken care of using normalization layers.

Our new notion of non-Euclidean convexity, the locally minimal perimeter, might have an appropriate generalization to graphs - which are also non-Euclidean. This probably involves extending our theory from parametric surfaces to Riemannian manifolds of general dimensions.

The representations we use for filtering may be used for other tasks, such as classification and segmentation - see preliminary result in Fig. 18.

Additional key aspects from the Euclidean case may be generalized to our parametric surface setting, e.g. curvature bounds of eigensets, and analysis of vectorial functions.

10 SUMMARY

We presented new nonlinear spectral theoretical analysis for surfaces, by generalizing nonlinear spectral theory of image processing. Based on our analysis, we proposed a general methodology for shape analysis and processing via nonlinear spectral filtering.

A key finding is our introduction of locally minimal perimeter sets, a novel generalization of conex sets to manifolds. It is derived by generalizing properties of *NETV* eigenfunctions. Our analysis is supported by numerical examples of minimizing flows, where numerical validation of eigenfunctions is performed by examining the decay near extinction, following the theory of [Bungert and Burger 2020].

For shape nonlinear filtering our methods extract spectral representations from smoothing flows which satisfy two requirements: zero-homogeneity and a fixed metric. We choose to process the

shape in its embedding space, providing unmediated nonlinear spectral representations, yielding good feature control. To showcase the general concept, three methods are proposed, where all three are based on the same mechanism, described in Eqs. (50), (53). Each method holds clear distinct properties induced by its flow, allowing various shape manipulations via spectral filtering. While possessing visibly distinct properties, all three methods demonstrate good smoothing and detail enhancement capabilities. Robustness to pose variations is demonstrated as well. With respect to processing time, we note that these methods are fairly fast, as they do not require solving an eigen-problem explicitly. Additionally, we present a Total-Variation approach for addressing the shape deformation problem. Our experiments show, that the deformation using our method is concentrated on plausible segment boundaries. Moreover - we have shown for several numerical cases, that these boundaries relate to our theoretical findings.¹²

REFERENCES

- Yonathan Aflalo, Ron Kimmel, and Dan Raviv. 2013. Scale invariant geometry for nonrigid shapes. *SIAM Journal on Imaging Sciences* 6, 3 (2013), 1579–1597.
- F. Andreu, C. Ballester, V. Caselles, and J. M. Mazón. 2001a. Minimizing total variation flow. *Differential and Integral Equations* 14, 3 (2001), 321–360.
- Fuensanta Andreu, Coloma Ballester, Vicent Caselles, and José M Mazón. 2001b. Minimizing total variation flow. (2001).
- Jean-François Aujol, Guy Gilboa, Tony Chan, and Stanley Osher. 2006. Structure-texture image decomposition—modeling, algorithms, and parameter selection. *International journal of computer vision* 67 (2006), 111–136.
- G. Bellettini, V. Caselles, and M. Novaga. 2002. The Total Variation Flow in \mathbb{R}^N . *Journal of Differential Equations* 184, 2 (2002), 475–525.
- Matania Ben-Artzi and Philippe G LeFloch. 2007. Well-posedness theory for geometry-compatible hyperbolic conservation laws on manifolds. In *Annales de l'Institut Henri Poincaré C, Analyse non linéaire*, Vol. 24. Elsevier, 989–1008.
- M. Benning, M. Möller, R.Z. Nossek, M. Burger, D. Cremers, G. Gilboa, and C.-B. Schönlieb. 2017. Nonlinear Spectral Image Fusion. In *Proceedings of the 6th International (SSVM'17)*. 41–53.
- Shai Biton and Guy Gilboa. 2022. Adaptive Anisotropic Total Variation: Analysis and Experimental Findings of Nonlinear Spectral Properties. *Journal of Mathematical Imaging and Vision* (2022), 1–23.
- Simon Blatt. 2009. A singular example for the Willmore flow. (2009).
- Mario Botsch and Olga Sorkine. 2007. On linear variational surface deformation methods. *IEEE transactions on visualization and computer graphics* 14, 1 (2007), 213–230.
- Amit Bracha, Oshri Halim, and Ron Kimmel. 2020. Shape correspondence by aligning scale-invariant LBO eigenfunctions. (2020).
- Alex Bronstein, Yoni Choukroun, Ron Kimmel, and Matan Sela. 2016. Consistent discretization and minimization of the l1 norm on manifolds. In *2016 Fourth International Conference on 3D Vision (3DV)*. IEEE, 435–440.
- Alexander M Bronstein, Michael M Bronstein, and Ron Kimmel. 2008. *Numerical geometry of non-rigid shapes*. Springer Science & Business Media.
- Leon Bungert and Martin Burger. 2020. Asymptotic profiles of nonlinear homogeneous evolution equations of gradient flow type. *Journal of Evolution Equations* 20, 3 (2020), 1061–1092.
- Leon Bungert, Martin Burger, Antonin Chambolle, and Matteo Novaga. 2019. Nonlinear spectral decompositions by gradient flows of one-homogeneous functionals. *Analysis & PDE* (2019).
- Martin Burger, Guy Gilboa, Michael Moeller, Lina Eckardt, and Daniel Cremers. 2016. Spectral decompositions using one-homogeneous functionals. *SIAM Journal on Imaging Sciences* 9, 3 (2016), 1374–1408.
- M. Burger, G. Gilboa, S. Osher, and J. Xu. 2006. Nonlinear inverse scale space methods. *Communications in Mathematical Sciences* 4, 1 (2006), 179–212.
- Martin Burger and Stanley Osher. 2013a. A guide to the TV zoo. In *Level set and PDE based reconstruction methods in imaging*. Springer, 1–70.
- Martin Burger and Stanley Osher. 2013b. A guide to the TV zoo. In *Level set and PDE based reconstruction methods in imaging*. Springer, 1–70.

¹²Models: Bust of Queen Nefertiti. Ägyptisches Museum und Papyrussammlung; Meteorite scan, courtesy of Jeremy Davis, Department of Art and Design, Central Michigan University; Trigon art; Stanford armadillo and poses by Belyaev, Yoshizawa, Seidel (2006); Michaels from [Bronstein et al. 2008]; various models from LIRIS database, Knubbel, Snail created by Thingiverse user 3D-mon, Cheburashka by Ilya Baran

- Simone Cammarasana and Giuseppe Patané. 2021. Localised and shape-aware functions for spectral geometry processing and shape analysis: A survey & perspectives. *Computers & Graphics* 97 (2021), 1–18.
- Antonin Chambolle. 2004. An algorithm for total variation minimization and applications. *Journal of Mathematical imaging and vision* 20, 1 (2004), 89–97.
- Antonin Chambolle, Vicent Caselles, Daniel Cremers, Matteo Novaga, and Thomas Pock. 2010. An introduction to total variation for image analysis. *Theoretical foundations and numerical methods for sparse recovery* 9, 263–340 (2010), 227.
- Antonin Chambolle and Thomas Pock. 2011. A first-order primal-dual algorithm for convex problems with applications to imaging. *Journal of mathematical imaging and vision* 40, 1 (2011), 120–145.
- Paolo Cignoni, Roberto Scopigno, and Marco Tarini. 2005. A simple normal enhancement technique for interactive non-photorealistic renderings. *Computers & Graphics* 29, 1 (2005), 125–133.
- Ido Cohen and Guy Gilboa. 2020. Introducing the p-Laplacian spectra. *Signal Processing* 167 (2020), 107281.
- Keenan Crane, Ulrich Pinkall, and Peter Schröder. 2013. Robust fairing via conformal curvature flow. *ACM Transactions on Graphics (TOG)* 32, 4 (2013), 1–10.
- Mathieu Desbrun, Mark Meyer, Peter Schröder, and Alan H Barr. 1999. Implicit fairing of irregular meshes using diffusion and curvature flow. In *Proceedings of the 26th annual conference on Computer graphics and interactive techniques*. 317–324.
- Julie Digne. 2012. Similarity based filtering of point clouds. In *2012 IEEE computer society conference on computer vision and pattern recognition workshops*. IEEE, 73–79.
- Chinthaka Dinesh, Gene Cheung, and Ivan V Bajić. 2019. 3D point cloud super-resolution via graph total variation on surface normals. In *2019 IEEE International Conference on Image Processing (ICIP)*. IEEE, 4390–4394.
- Chinthaka Dinesh, Gene Cheung, and Ivan V Bajić. 2020. Super-resolution of 3D color point clouds via fast graph total variation. In *ICASSP 2020-2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 1983–1987.
- Manfredo P Do Carmo. 2016. *Differential geometry of curves and surfaces: revised and updated second edition*. Courier Dover Publications.
- Puhong Duan, Xudong Kang, Shutao Li, and Pedram Ghamisi. 2019. Noise-robust hyperspectral image classification via multi-scale total variation. *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing* 12, 6 (2019), 1948–1962.
- Abderrahim Elmoataz, Olivier Lezoray, and Sébastien Bouteux. 2008. Nonlocal discrete regularization on weighted graphs: a framework for image and manifold processing. *IEEE transactions on Image Processing* 17, 7 (2008), 1047–1060.
- Marco Fumero, Michael Möller, and Emanuele Rodolà. 2020. Nonlinear spectral geometry processing via the tv transform. *ACM Transactions on Graphics (TOG)* 39, 6 (2020), 1–16.
- Sylvestre Gallot, Dominique Hulin, and Jacques Lafontaine. 1990. *Riemannian geometry*. Vol. 2. Springer.
- Sylvestre Gallot, Dominique Hulin, and Jacques Lafontaine. 2004. *Differential Manifolds*. In *Riemannian Geometry*. Springer, 1–49.
- Guy Gilboa. 2013. A spectral approach to total variation. In *Int. Conf. on Scale Space and Variational Methods in Computer Vision*. Springer, 36–47.
- Guy Gilboa. 2014a. A total variation spectral framework for scale and texture analysis. *SIAM journal on Imaging Sciences* 7, 4 (2014), 1937–1961.
- G. Gilboa. 2014b. A total variation spectral framework for scale and texture analysis. *SIAM Journal on Imaging Sciences* 7, 4 (2014), 1937–1961.
- Guy Gilboa. 2018. *Nonlinear Eigenproblems in Image Processing and Computer Vision*. Springer.
- Markus Grasmair and Frank Lenzen. 2010. Anisotropic total variation filtering. *Applied Mathematics & Optimization* 62, 3 (2010), 323–339.
- Batu Güneysu and Diego Pallara. 2015. Functions with bounded variation on a class of Riemannian manifolds with Ricci curvature unbounded from below. *Math. Ann.* 363, 3 (2015), 1307–1331.
- Ester Hait and Guy Gilboa. 2019. Spectral total-variation local scale signatures for image manipulation and fusion. *IEEE Trans. Image Processing* 28, 2 (2019), 880–895.
- Gerhard Huisken. 1990. Asymptotic behavior for singularities of the mean curvature flow. *Journal of Differential Geometry* 31, 1 (1990), 285–299.
- Michael Kazhdan, Jake Solomon, and Mirela Ben-Chen. 2012. Can mean-curvature flow be modified to be non-singular?. In *Computer Graphics Forum*, Vol. 31. Wiley Online Library, 1745–1754.
- Bertrand Kerautret and Jacques-Olivier Lachaud. 2020. Geometric Total Variation for Image Vectorization, Zooming and Pixel Art Depixelizing. In *Pattern Recognition: 5th Asian Conference, ACPR 2019, Auckland, New Zealand, November 26–29, 2019, Revised Selected Papers, Part I 5*. Springer, 391–405.
- Kimmel, Malladi, and Sochen. 2000. Images as embedded maps and minimal surfaces: movies, color, texture, and volumetric medical images. *International Journal of Computer Vision* 39, 2 (2000), 111–129.
- Ron Kimmel, Ravi Malladi, and N Sochen. 1998. Image processing via the Beltrami operator. In *Asian Conference on Computer Vision*. Springer, 574–581.
- John M Lee. 2013. Smooth manifolds. In *Introduction to Smooth Manifolds*. Springer, 1–31.
- Chengcai Leng, Hai Zhang, Guorong Cai, Zhen Chen, and Anup Basu. 2021. Total variation constrained non-negative matrix factorization for medical image registration. *IEEE/CAA Journal of Automatica Sinica* 8, 5 (2021), 1025–1037.
- Or Litany, Tal Remez, Emanuele Rodola, Alex Bronstein, and Michael Bronstein. 2017. Deep functional maps: Structured prediction for dense shape correspondence. In *Proceedings of the IEEE international conference on computer vision*. 5659–5667.
- Hsueh-Ti Derek Liu, Alec Jacobson, and Keenan Crane. 2017. A Dirac operator for extrinsic shape analysis. In *Computer Graphics Forum*, Vol. 36. Wiley Online Library, 139–149.
- Michele Miranda Jr. 2003. Functions of bounded variation on “good” metric spaces. *Journal de mathématiques pures et appliquées* 82, 8 (2003), 975–1004.
- Neal Parikh, Stephen Boyd, et al. 2014. Proximal algorithms. *Foundations and trends® in Optimization* 1, 3 (2014), 127–239.
- Barbara Pascal, Samuel Vaiter, Nelly Pustelnik, and Patrice Abry. 2021. Automated data-driven selection of the hyperparameters for total-variation-based texture segmentation. *Journal of Mathematical Imaging and Vision* 63, 7 (2021), 923–952.
- L. Rudin, S. Osher, and E. Fatemi. 1992. Nonlinear total variation based noise removal algorithms. *Physica D* 60 (1992), 259–268.
- Shrutika S Sawant and Manoharan Prabukumar. 2020. A review on graph-based semi-supervised learning methods for hyperspectral image classification. *The Egyptian Journal of Remote Sensing and Space Science* 23, 2 (2020), 243–248.
- Matan Sela, Yonathan Aflalo, and Ron Kimmel. 2015. Computational caricaturization of surfaces. *Computer Vision and Image Understanding* 141 (2015), 1–17.
- Nicholas Sharp and Keenan Crane. 2020. A laplacian for nonmanifold triangle meshes. In *Computer Graphics Forum*, Vol. 39. Wiley Online Library, 69–80.
- Olga Sorkine and Marc Alexa. 2007. As-rigid-as-possible surface modeling. In *Symposium on Geometry processing*, Vol. 4. 109–116.
- Olga Sorkine and Mario Botsch. 2009. Interactive Shape Modeling and Deformation.. In *Eurographics (Tutorials)*. 11–37.
- Olga Sorkine, Daniel Cohen-Or, Yaron Lipman, Marc Alexa, Christian Rössl, and H-P Seidel. 2004. Laplacian surface editing. In *Proceedings of the 2004 Eurographics/ACM SIGGRAPH symposium on Geometry processing*. 175–184.
- Michael Spivak. 2018. *Calculus on manifolds: a modern approach to classical theorems of advanced calculus*. CRC press.
- Gabriel Taubin. 1995. A signal processing approach to fair surface design. In *Proceedings of 22nd annual conf. Computer graphics and techniques*. 351–358.
- Bruno Vallet and Bruno Lévy. 2008. Spectral geometry processing with manifold harmonics. In *Computer Graphics Forum*, Vol. 27. Wiley Online Library, 251–260.
- Max Wardetzky, Saurabh Mathur, Felix Kälberer, and Eitan Grinspun. 2007. Discrete Laplace operators: no free lunch. In *Symposium on Geometry processing*. Aire-la-Ville, Switzerland, 33–37.
- Wei Wei, Bin Zhou, Dawid Polap, and Marcin Woźniak. 2019. A regional adaptive variational PDE model for computed tomography image reconstruction. *Pattern Recognition* 92 (2019), 64–81.
- Wetzler, Aflalo, Dubrovina, and Kimmel. 2013. The Laplace-Beltrami operator: a ubiquitous tool for image and shape processing. In *Int. Symp. on Mathematical Morphology and Applications to Signal and Image Processing*. Springer, 302–316.
- Wang Yifan, Lukas Rahmann, and Olga Sorkine-Hornung. 2021. Geometry-consistent neural shape representation with implicit displacement fields. *arXiv preprint arXiv:2106.05187* (2021).
- Huayan Zhang and Chunxue Wang. 2020. Total variation diffusion and its application in shape decomposition. *Computers & Graphics* 90 (2020), 95–107.
- Huayan Zhang, Chunlin Wu, Juyong Zhang, and Jiansong Deng. 2015. Variational mesh denoising using total variation and piecewise constant function space. *IEEE transactions on visualization and computer graphics* 21, 7 (2015), 873–886.
- Jianwei Zhang, Jing Qi, Zhaohui Zheng, and Le Sun. 2022. A robust image segmentation framework based on total variation spectral transform. *Pattern Recognition Letters* 153 (2022), 159–167.
- Saishang Zhong, Zhong Xie, Weina Wang, Zheng Liu, and Ligang Liu. 2018. Mesh denoising via total variation and weighted Laplacian regularizations. *Computer Animation and Virtual Worlds* 29, 3–4 (2018), e1827.

A COMPLEMENTARY EXPERIMENTS

Here we add experiments for completeness. Fig. 20 bridges a gap between Figs. 3 and 6. Fig. 21 shows shape exaggeration similar to Fig. 12 on additional poses while not using smoothing. Fig. 19 extends the method compariso performed for the Snail model in Fig. 15 to three other models.

B SOME PROOFS REGARDING SEC. 4.1

CLAIM. *Let a field z that is normal to the boundary of C on (almost all) boundary points, and of norm less than or equal to one everywhere*




Fig. 19. The same experiment of Fig. 15, for the other shapes shown in Fig. 13. As before the shape and deformation constraints are tested for the three methods. The same properties observed in Fig. 15 can be seen here as well: The deformation field is piece-wise constant, and the deformation is concentrated on boundaries of plausible shape segments.




Fig. 20. A similar experiment to the one shown in Fig. 3, comparing linear diffusion (upper row) to the NETV minimizing flow (bottom row). This experiment is performed using the z-coordinate function. The resulting flow was used to induce the spectral representations filtered in Fig. 6.

on M , i.e.

$$z = \tilde{n}^C \text{ for a.e. } \omega_1, \omega_2 \in \partial\tilde{C}, \quad ||z||_g \leq 1 \quad \forall \omega_1, \omega_2 \in \Omega, \quad (59)$$

where a.e. stands for "almost every". Then z admits the supremum of Eq. (24) for $NETV(C)$ ¹³.

PROOF. Since $\tilde{\chi}^C$ is an indicator of \tilde{C} we have

$$\int_{\Omega} \nabla_g \cdot z \tilde{\chi}^C da = \int_{\tilde{C}} \nabla_g \cdot z da. \quad (60)$$

¹³Requiring this almost everywhere on $\partial\tilde{C}$ is enough, as our proofs use z under integration.



Fig. 21. Band-Pass exaggeration via $M3$ for the armadillo in various poses, where all poses experience the same filter. The band-pass exaggeration magnitude is identical to Fig. 12. Unlike Fig. 12, there is no smoothing of the high frequencies. The consistency across poses shows that our method can be considered as a caricaturization method by the properties posed in [Sela et al. 2015].

Using manifold divergence Thm., as stated in Eq.(10), we have

$$\int_{\Omega} \nabla_g \cdot z \tilde{\chi}^C da = \int_{t_1}^{t_2} \langle z, \tilde{n}^C \rangle_g ||\tilde{\gamma}_t^C||_g dt, \quad (61)$$

where $\tilde{\gamma}_t^C = \frac{d\tilde{\gamma}^C}{dt}$. For convenience - let us reformulate this as

$$\int_{\Omega} \nabla_g \cdot z \tilde{\chi}^C da = \int_{t_1}^{t_2} ||z||_g \cos\theta ||\tilde{\gamma}_t^C||_g dt, \quad (62)$$

where $\theta = \angle_g(z, \tilde{n}^C) = \cos^{-1} \frac{\langle \tilde{n}^C, z \rangle_g}{||\tilde{n}^C||_g ||z||_g}$. Maximization under the constraint $||z||_g \leq 1$ is achieved for $\theta = 0$, and $||z||_g = 1$ almost everywhere on the boundary, i.e.

$$z = \tilde{n}^C \text{ for a.e. } \omega_1, \omega_2 \in \partial\tilde{C} \Rightarrow NETV(C) = \int_{\Omega} \nabla_g \cdot z \tilde{\chi}^C da. \quad (63)$$

□

Remark: The "almost all" condition allows robustness to a zero-measure subset of $\partial\tilde{C}$ in which boundary normals are not defined (namely points of non-differentiable $\partial\tilde{C}$). In such a case \tilde{n}^C may be extended to satisfy

$$\begin{cases} J\tilde{n}^C = n^C \quad \forall \omega_1, \omega_2 \in \partial\tilde{C} & \text{if } n^C \text{ exists} \\ ||\tilde{n}^C||_g \leq 1 & \text{otherwise,} \end{cases} \quad (64)$$

while keeping the proof intact.

CLAIM.

$$NETV(C) = per(C). \quad (65)$$

This claim can be found for instance in [Fumero et al. 2020]. Let us re-prove it in our setting:

PROOF. By Eq. (61) we have

$$NETV(C) = \sup_z \int_{t_1}^{t_2} \langle z, \tilde{n}^C \rangle_g \|\tilde{y}_t^C\|_g dt \text{ s.t. } \|z\|_g \leq 1 \forall \omega_1, \omega_2 \in \Omega. \quad (66)$$

By Eq. (63) we can obtain a supremum by assigning $z = \tilde{n}^C \forall \omega_1, \omega_2 \in \partial\tilde{C}$ and have

$$NETV(C) = \int_{t_1}^{t_2} \|\tilde{n}^C\|_g^2 \|\tilde{y}_t^C\|_g dt = \int_{t_1}^{t_2} \|\tilde{y}_t^C\|_g dt = per(C), \quad (67)$$

where the second equality uses $\|\tilde{n}^C\|_g = 1 \forall \omega_1, \omega_2$, and the last equality comes from Def. (5). \square

C SLEEVE SETS AS EIGENSETS ON THE TORUS

Here we show that sleeve sets of the torus are eigensets, considering M to be a torus with big and small radii R, r .

C.1 Torus preliminaries

We choose its parametric formulation S as follows: $S(\omega_1, \omega_2) = ((R + r \cos(\omega_1)) \cos(\omega_2), (R + r \cos(\omega_1)) \sin(\omega_2), r \sin(\omega_1))^T$, where $\omega_1, \omega_2 \in \Omega$ and $\Omega = [-\pi, \pi) \times [-\pi, \pi)$, inducing a metric $g(\omega_1, \omega_2) = \begin{pmatrix} (R + r \cos(\omega_1))^2 & 0 \\ 0 & r^2 \end{pmatrix}$, resulting with

$$\sqrt{|g|} = (R + r \cos(\omega_1)) r. \quad (68)$$

Let z be a field on the torus, then its squared norm function is

$$\|z\|_g = \sqrt{(R + r \cos(\omega_1))^2 z[1]^2 + r^2 z[2]^2}, \quad (69)$$

where $z[1], z[2]$ are components of the field. Remembering the divergence formula

$$\nabla_g \cdot \tilde{F} = \frac{1}{\sqrt{|g|}} \nabla_{\omega_1, \omega_2} \cdot (\sqrt{|g|} \tilde{F}), \quad (70)$$

(see for instance [Do Carmo 2016]) we have for the torus:

$$\nabla_g \cdot z = \frac{1}{(R + r \cos(\omega_1)) r} \nabla_{\omega_1, \omega_2} \cdot [(R + r \cos(\omega_1)) \nabla z] \quad (71)$$

$$\nabla_g \cdot z = \frac{\partial}{\partial \omega_1} z[1] + \frac{-r \sin(\omega_1)}{R + r \cos(\omega_1)} z[1] + \frac{\partial}{\partial \omega_2} z[2]. \quad (72)$$

The sleeve set of angle-length l , and center at $\omega_2 = c_0$ is denoted in parameterization domain as \tilde{C} , and defined as follows:

$$\tilde{C} = \{\omega_1, \omega_2 : |\omega_2 - c_0| \leq \frac{l}{2}\}. \quad (73)$$

For convenience, W.T.L.O.G. we consider the sleeve set to have center at $\omega_2 = 0$, i.e.

$$\tilde{C} = \{\omega_1, \omega_2 : |\omega_2| \leq \frac{l}{2}\}, \quad (74)$$

resulting with

$$\partial\tilde{C} = \{\omega_1, \omega_2 : |\omega_2| = \frac{l}{2}\}. \quad (75)$$

C.2 Finding a field of required properties

The first property we are looking for is orthogonality to the boundary, on all boundary points. Since g is diagonal, we have that S preserves angles - hence orthogonality may be tested in parameterization domain. Furthermore, we need the field to be of unit norm on the boundary. By (75) we have that z is orthogonal to the boundary and of unit norm on the boundary if $z(\omega_1, \omega_2 = \frac{l}{2})[1] = \frac{1}{R + r \cos(\omega_1)}$, $z(\omega_1, \omega_2 = -\frac{l}{2})[1] = -\frac{1}{R + r \cos(\omega_1)}$, and $z(\omega_1, \omega_2 = \frac{l}{2})[2] = 0 \forall \omega_1$.

$$\text{Let } \Theta(\omega_2) = \begin{cases} \omega_2 - \frac{l}{2} & \omega_2 \in [0, \pi) \\ \omega_2 + \frac{l}{2} & \omega_2 \in (-\pi, 0) \end{cases}. \text{ The choice}$$

$$z = \left(\frac{\frac{2}{l} \omega_2}{R + r \cos(\omega_1)}, \alpha \Theta(\omega_2) \right)^T, \quad (76)$$

satisfies above properties. Another required property is a constant divergence inside C . Let us show that this choice satisfies that as well, except for a zero-measure set of points: Plugging to (72) we have

$$\nabla_g \cdot z = -\frac{2}{l} \omega_2 \frac{-r \sin(\omega_1)}{(R + r \cos(\omega_1))^2} + \frac{-r \sin(\omega_1)}{R + r \cos(\omega_1)} \frac{\frac{2}{l} \omega_2}{R + r \cos(\omega_1)} + \alpha \frac{\partial \Theta(\omega_2)}{\partial \omega_2}, \quad (77)$$

where $\frac{\partial \Theta(\omega_2)}{\partial \omega_2} = 1 \forall \omega_2 \neq 0, \pi$, thus

$$\nabla_g \cdot z = \alpha \forall \omega_2 \neq 0, \pi. \quad (78)$$

C.3 Demonstrating eigensets

To prove an eigenset, we need to construct a field ξ , that satisfies the eigenfunction properties (5.1) for a ψ as in (27). In the current case ψ is defined for a sleeve set C on a torus.

First we note, that the appropriate field z of (76) can be similarly defined for the complement set - which is a sleeve set as well, but with its center at $\omega_2 = \pi$. It turns out, that the fields for center at either $\omega_2 = \pi$, or at $\omega_2 = 0$, are of the same form, up to a sign factor. Thus we define ξ to admit (76) for \tilde{C} inside \tilde{C} , and the minus of (76) for $\Omega \setminus \tilde{C}$, i.e.

$$\xi = \begin{cases} \left(\frac{\frac{2}{l} \omega_2}{R + r \cos(\omega_1)}, \alpha_1 \Theta(\omega_2) \right)^T & \omega_1, \omega_2 \in \tilde{C} \\ \left(\frac{\frac{2}{l} \omega_2}{R + r \cos(\omega_1)}, \alpha_2 \Theta(\omega_2) \right)^T & \omega_1, \omega_2 \in \Omega \setminus \tilde{C} \end{cases}, \quad (79)$$

which has $\nabla_g \cdot \xi = \begin{cases} \alpha_1 & \omega_1, \omega_2 \in \tilde{C} \\ \alpha_2 & \omega_1, \omega_2 \in \Omega \setminus \tilde{C} \end{cases}$, and is unit-orthogonal to

the boundary, i.e. $z = \tilde{n}^C \forall \omega_1, \omega_2 \in \partial\tilde{C}$. If we set $\alpha_1 = 1, \alpha_2 = \beta$, where β is as in (28), then indeed we have the eigenfunction property $\psi = \lambda \nabla_g \cdot \xi$.

By (5.1) it is left to show that $\nabla_g \cdot \xi \in \partial NETV(\psi)$. To show this, it is sufficient, by Eq. (19), to show that $\int_{\Omega} \nabla_g \cdot \xi \psi da = NETV(\psi)$. Let us begin: By one-homogeneity of the $NETV$ - we have that $NETV(\psi) = (1 + \beta) NETV(C)$.

Since $\xi = \tilde{n}^C \forall \omega_1, \omega_2 \in \partial\tilde{C}$, we have by Eq. (63) that $NETV(C) = \int_{\Omega} \nabla_g \cdot \xi \tilde{\chi}^C da$.

Thus by the definition of ψ (27), we have

$$\int_{\Omega} \nabla_g \cdot \xi \psi da = \int_{\Omega} \nabla_g \cdot \xi (\tilde{\chi}^C - \beta \tilde{\chi}^{M \setminus C}) da =$$

$$\begin{aligned}
& \int_{\Omega} \nabla g \cdot \xi (\tilde{\chi}^C - \beta(1 - \tilde{\chi}^C)) da = \\
& (1 + \beta) \int_{\Omega} \nabla g \cdot \xi \tilde{\chi}^C da - \beta \int_{\Omega} \nabla g \cdot \xi da = \\
& (1 + \beta) \int_{\Omega} \nabla g \cdot \xi \tilde{\chi}^C da - \lambda \beta \int_{\Omega} \psi d\bar{a} = \\
& (1 + \beta) \int_{\Omega} \nabla g \cdot \xi \tilde{\chi}^C da,
\end{aligned}$$

where the last cancellation uses the Neumann boundary condition assumption. Thus we have $\int_{\Omega} \nabla g \cdot \xi \psi da = (1 + \beta) \int_{\Omega} \nabla g \cdot \xi \tilde{\chi}^C da = (1 + \beta) \text{NETV}(C) = \text{NETV}(\psi)$. By Eq. (19) we know that this is sufficient for $\xi \in \partial \text{NETV}(\psi)$.

D SHAPE DEFORMATION DETAILS

First, let us consider for simplicity the scalar constrained *NETV* flow: Consider a manifold M and a function $f(q)$ on the manifold. Let us perform a total variation minimizing flow, initialized with f . The evolving function is constrained throughout the flow, i.e. denote $u(q, t)$ as the function at time t of the flow, and the constraints $\{u(q_i, t) = c(q_i), q_i \in M\}_{i=1}^n$. Denote $u(q) := u(q, t = \infty)$.

Similarly, we may assume a surface M , with x as its coordinate functions. The shape deformation problem requires to find some new surface with coordinates x' , which preserves the natural structure of x , under the constraints that some points are pre-determined. To this end we define the deformation field of a proposed x' as $u = x' - x$ (from which x' is reconstructed as $x' = x + u$). We perform constrained total-variation minimization on u , where each nonzero constraint in $\{c(q_i)\}_{i=1}^n$ is translated to a point $x'(q_i) \neq x(q_i)$. For the minimization process, we adapt a vectorial version of the process introduced by [Bronstein et al. 2016], where we use p_{M1} of (54) as the sub-differential. See Figs. 15, 13, 19 for the results.

E SHAPE FLOWS AND THEIR EQUIVALENCE SETS

This section serves as a reminder of observations on equivalent flows, namely the equivalence presented in [Kimmel et al. 1998]. Consider the flow equation,

$$\frac{\partial f}{\partial t} = p(f(t)), \quad (80)$$

where p is some operator. E.g., choosing $p = \Delta_{g,2}$ we obtain Eq. (12). Suppose that M is a manifold of some shape, and we would like to process this shape via a flow. To do so we can initialize the flow with the shape's coordinate function, i.e. $f(0) = S$. Doing so with Eq. (12) is a classical shape smoothing flow. During the shape flow, at a given time t , we have $p(f(t)) : M \rightarrow \mathbb{R}^3$, a field which describes the velocity of the evolving shape's points. Recalling the normal

$$N = \frac{S_{\omega_1} \times S_{\omega_2}}{|S_{\omega_1} \times S_{\omega_2}|}, \quad (81)$$

a vector $\vec{V} \in p(f(t))$ may be decomposed to its normal and tangential component as follows: $\vec{V}_N = \langle \vec{V}, \vec{N} \rangle \vec{N}$, $\vec{V}_T = \vec{V} - \vec{V}_N$, yielding $\frac{\partial f}{\partial t} = p(f)_N + p(f)_T$. The normal component accounts for the change of the shape in time, and the tangential movement is merely a change of the shape's parameterization in time. Thus two shape flows are considered equal if their normal components are equal, and an equivalence set of shape flows is defined as:

$$\left\{ \frac{\partial f}{\partial t} = q(f) : \langle q(f), \vec{N} \rangle = \langle p(f), \vec{N} \rangle \right\}. \quad (82)$$

We note that parameterization constraints may cause differences between two flows from the same equivalence set.

E.1 Beltrami -TV Flow equivalence

In [Kimmel et al. 1998], an equivalence between MCF and the TV-flow was shown: On one hand they define the Beltrami flow $\frac{\partial f}{\partial t} = \frac{H(t)}{\langle \vec{N}(t), \hat{z} \rangle} \hat{z}$, which is obviously equivalent to MCF, in the sense of Eq. (82). On the other hand, consider a shape parameterized as an "image function", i.e. $S = (u, v, f(u, v))$, where the image is given by $f(u, v)$, and u, v are a 2D Euclidean domain discretized as the pixel grid. In this case they show that plugging Eq. (81) in to the Beltrami flow is equivalent to the image TV-flow.

Remark: equivalence by Eq. (82) does not account for parameterization, which may induce implicit constraints - as is the case here: During Beltrami flow, the evolving shape's points are constrained to move in the \hat{z} direction, thus keeping the parameterization $S = (u, v, f(u, v))$, contrary to MCF, where no such constraint exists. This is the reason MCF converges to a point, while image TV flow converges to a plane (u, v, const) .