



Nonlinear Spectral Processing of Shapes via Zero-Homogeneous Flows

Jonathan Brokman^(✉) and Guy Gilboa^(✉)

Technion - Israel Institute of Technology, 3200003 Haifa, Israel
sjonatha@campus.technion.ac.il, guy.gilboa@ee.technion.ac.il

Abstract. In this work we extend the spectral total-variation framework, and use it to analyze and process 2D manifolds embedded in 3D. Analysis is performed in the embedding space - thus “spectral arithmetics” manipulate the shape directly. This makes our approach highly versatile and accurate for feature control. We propose three such methods, based on non-Euclidean zero-homogeneous p -Laplace operators. Each method satisfies distinct characteristics, demonstrated through smoothing, enhancing and exaggerating filters.

Keywords: 1-Laplacian · 3-Laplacian · Spectral TV · cMCF · Mesh processing · Nonlinear spectral processing

1 Introduction

In 1995 Taubin [22] 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. [21]). Over time, the Laplace-Beltrami became the standard Laplacian of choice, for spectral applications, and in general [24]¹.

Nonlinear spectral processing has been developed in recent years for image analysis and manipulation. Spectral total-variation (TV) was introduced in [13, 14] facilitating nonlinear edge-preserving filtering. Essentially, the idea is to decompose a signal into nonlinear spectral elements related to eigenfunctions of the total-variation subgradient. The method is based on evolving gradient descent with respect to the TV functional (TV-flow [1]). The spectral elements decay linearly in this flow. Different decay rates correspond to different scales

¹ Such an adaptation of [22] can be found e.g. in [23], which proposed a computationally efficient shape filtering, and demonstrated some core filtering capabilities: Shape exaggeration, detail enhancement, shape smoothing and regularization.

We acknowledge support by grant agreement No. 777826 (NoMADS), by the Israel Science Foundation (Grant No. 534/19) and by the Ollendorff Minerva Center.

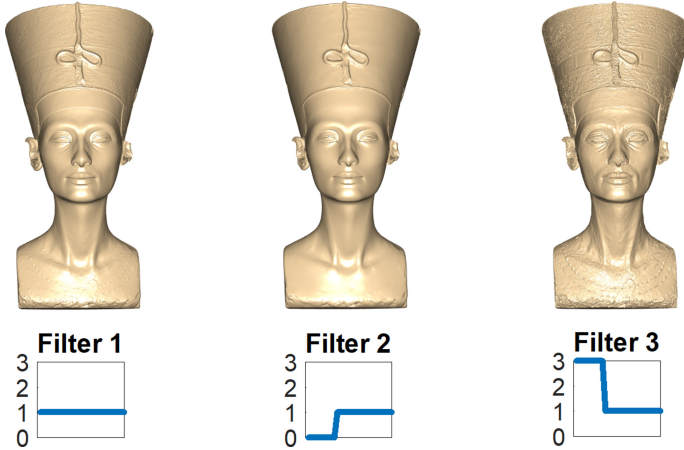


Fig. 1. Bust of Queen Nefertity reconstructed and filtered via Directional Shape TV (M3)

(in the case of a single eigenfunction the rate is exactly the eigenvalue). Theoretical underpinning was performed for the spatial discrete case [7] and continuous case [5]. 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 [2, 15]. Only recently, in 2020, a spectral TV framework for shape processing was proposed, for the first time, by Fumero et al. [12]. They advocate applying spectral TV to the normals of the shape, thus gradients are calculated on the normals' domain (unit sphere), and spectral processing of the embedded shape is done implicitly.

Burger et al. [7] generalized 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 [6] were also proposed. The space-continuous setting was recently analyzed in [5]. A common thread related to gradient flows of one-homogeneous functionals is that they are based on zero-homogeneous operators.

In this work we propose three new methods, extending nonlinear spectral processing of shapes using zero-homogeneous flows. We modify existing Laplacian-based flows of shapes such that they will comply with the prerequisites of nonlinear spectral processing. We demonstrate how each flow induces distinct spectral properties. On one hand, our work extends ideas from [12] to other zero-homogeneous flows. On the other hand - we examine a complementary approach, as our proposed spectral processing methods are performed explicitly on the shape in its embedding space.

2 Background

2.1 Differential Operators on Manifolds

Continuous Formulation. In this work, the processed shape is assumed to be a smooth 2D manifold embedded in 3D, denoted by M , represented as $S(\omega_1, \omega_2) = (x(\omega_1, \omega_2), y(\omega_1, \omega_2), z(\omega_1, \omega_2))$, i.e. $S : \Omega \subset \mathbb{R}^2 \rightarrow M$. Let $f : S \rightarrow \mathbb{R}$, and $\tilde{f} = f \circ S(\omega_1, \omega_2)$ i.e. $\tilde{f} : \Omega \rightarrow \mathbb{R}$. For instance, we can map each point $q = S(\omega_1, \omega_2) \in M$ to $x(\omega_1, \omega_2)$. This function is termed the *x-coordinate function*. If f is the x-coordinate function, then \tilde{f} maps each point ω_1, ω_2 to $x(\omega_1, \omega_2)$. Let $T_q M$ be the plane tangent to M at point $q \in M$, it can be shown that $\frac{\partial S}{\partial \omega_1}, \frac{\partial S}{\partial \omega_2} \in \mathbb{R}^3$, denoted $S_{\omega_1}, S_{\omega_2}$, at point q span $T_q M$. M is equipped with a metric g ,

$$g(\omega_1, \omega_2) = \begin{pmatrix} S_{\omega_1}^T S_{\omega_1} & S_{\omega_1}^T S_{\omega_2} \\ S_{\omega_2}^T S_{\omega_1} & S_{\omega_2}^T S_{\omega_2} \end{pmatrix}, \quad (1)$$

which, given that $S_{\omega_1}, S_{\omega_2}$ are linearly independent, can be shown to be positive semi-definite and invertible. g induces the inner product $\langle a, b \rangle_g = a^T g b = A^T B$, where $A, B \in T_q M$ are the mapping of $a, b \in \Omega$, considering mapped vectors to be velocities of mapped routes. The squared-root determinant satisfies the ‘‘area elements’’ $\sqrt{|g|} du dv$. The gradient operator $\nabla_g f(q) = g^{-1} \nabla_{\omega_1, \omega_2} \tilde{f}(\omega_1, \omega_2)$ satisfies $\langle \nabla_g f(q), \mathbf{w} \rangle = \lim_{h \rightarrow 0} \frac{f(q+h\mathbf{w}) - f(q)}{h}$, and the divergence $\nabla_g \cdot \mathbf{F} = \frac{1}{\sqrt{|g|}} \nabla_{\omega_1, \omega_2} \cdot (\sqrt{|g|} \tilde{\mathbf{F}})$ is obtained as an adjoint of $\nabla_g(f)$. Full details are available at [10]. Finally, the \mathcal{P} -Laplace-Beltrami is defined by,

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

For $\mathcal{P} = 2$ the Laplace-Beltrami is obtained. Other special cases we will discuss are $\mathcal{P} = 3$, and $\mathcal{P} = 1$.

Discretization. We use the discretization framework of [16]. M is approximated as a triangulated mesh and f as a function on vertices. Discrete matrix operators are derived, including $[\sqrt{|g|} \cdot]$ as mass matrix denoted $[M]$, $[\nabla_g]$ denoted $[G]$ and $[\sqrt{|g|} \cdot \nabla_g]$ denoted $[D]$. Thus $\nabla_g \cdot$ is discretized as $[M^{-1}][D]$. We remark that both $[G]$ and $[M^{-1}][D]$ satisfy linearity, are Hermitian conjugate of one another, and $[M^{-1}][D][G]$ is a common discretization of the Laplace-Beltrami operator. Finally, a semi-discrete diffusion process of f on M can be formulated as,

$$\frac{\partial u(t)}{\partial t} = [M^{-1}][D][G] \cdot u(t), \quad u(0) = f. \quad (3)$$

2.2 Vectorial Total Variation

Total variation (TV) has been used extensively for the past three decades in image processing and computer vision (see [8,9] for theory and applications). For smooth functions, the TV functional over the domain $\Omega \subset \mathbb{R}^n$ is $TV(u) = \int_{\Omega} |\nabla u(x)| dx$, where $x = \{x_1, \dots, x_n\}$. When dealing with a vectorial function (of several channels) it is desired to take into account the

inter-component correlations. A common definition for vectorial TV (see e.g. [3]) is: $VTV(u) = \int_{\Omega} \sqrt{\sum_c |\nabla u_c(x)|^2} dx$, where u_c is channel $c = 1, \dots, C$ in $u = \{u_1, \dots, u_C\}$. The gradient descent flow evolves each channel by the VTV flow: $\frac{\partial u_c}{\partial t} = \nabla \cdot \left(\frac{\nabla u_c}{\sqrt{\sum_{\varepsilon} |\nabla u_{\varepsilon}|^2}} \right)$. Note, as in the scalar TV-flow, this flow is zero-homogeneous.

2.3 Laplacian-Based Flows

Mesh smoothing is often achieved by some form of a discrete diffusion process. A plethora of such algorithms were proposed, here we present the three most relevant to this paper.

Heat Flow. The simplest of the three, often termed *Heat Flow*, smooths each coordinate function via Eq. (3), where $[M], [D], [G]$ are induced by the initial shape’s metric, denoted g_0 and are fixed throughout the process. The smoothed shape is given by the final three smoothed coordinate functions.

Mean Curvature Flow. The most notable shape smoothing process, *Mean Curvature Flow* (MCF) is derived as an area minimizing process. A common implementation of it resembles the heat flow method, however, at each time step the operator matrices $[M], [D], [G]$ are re-calculated according to the present metric of the smoothed shape, denoted g_t [19].

Conformalized Mean Curvature Flow. In [17] cMCF was introduced as a modified version of MCF. The metric is updated at each time step, but unlike MCF, the metric is conformalized to the initial shape’s metric. This conformalized metric, $\tilde{g}_t = \sqrt{|g_0^{-1} g_t|} g_0$, induces the “conformalized Laplace-Beltrami”. Proposed implementation is similar to the former whereas $[M]$ is updated w.r.t g_t while $[D], [G]$ are fixed w.r.t g_0 . cMCF is significantly more immune to singularities than MCF, making it more suitable for editing surface extremities such as head and limbs of humanoid models. Adding rescaling for numerical stability, the shape converges to a sphere (not a point). In [17] they demonstrate that the resulting smoothed shape admits a conformal mapping of the initial shape.

2.4 TV Mesh Processing

A suitable TV flow for M is a nontrivial task, as its immediate representation $S(u, v) = x(u, v), y(u, v), z(u, v)$ is a vectorial function of correlated components. We briefly summarize two methods relevant to our work, related to graph smoothing and to spectral processing of normals.

\mathcal{P} -Laplace Flows on Weighted Graphs. Elmoataz et al. [11] proposed a generalization of TV regularization of functions on weighted graphs. Their discrete gradient and divergence operators, induced by the graph topology, provide a non-Euclidean \mathcal{P} -Laplace. Treating mesh triangle sides as graph edges, \mathcal{P} -Laplace mesh smoothing was proposed. Shape’s inter correlations were accounted via combined gradient magnitude as in VTV.

Normal TV: Flow and Shape Spectral Analysis. Recently, in [12] a spectral TV framework for shape analysis was suggested. The authors proposed spectral TV analysis of the shape’s normals, followed by shape-from-normals reconstruction. This achieves rotation invariance while operating on a fixed metric (the unit-sphere). Another interesting property of this implicit approach is the convergence to a plain, which is a translation of normals converging to a point. They show the method is compatible with a range of applications.

3 Proposed Methods

We propose a general framework for nonlinear spectral filtering of shapes (meshes). First, we present our zero-homogeneous framework. Then we suggest three methods which utilize this framework to filter shapes. Each method is inspired by a different flow: Heat Flow, cMCF and MCF. The shape is represented in its embedding space, allowing excellent feature control as spectral manipulations directly manipulate embedded structures.

Spectral Processing via (Any) Zero-Homogeneous Flow. Let X be a space of functions on a general Euclidean space. Let $p : X \rightarrow X$ be a zero-homogeneous operator, i.e.,

$$p(\alpha f) = \text{sign}(\alpha)p(f), \quad \alpha \in \mathbb{R}, f \in X. \quad (4)$$

We examine the following zero-homogeneous flow,

$$u_t = -p(u(t)), \quad u(0) = f_0 \in X, \quad t \geq 0, \quad (5)$$

where $u_t = \frac{\partial u}{\partial t}$. We assume the flow exists and that the solution is unique. We also assume the second time-derivative of u exists (in the distributional sense) a.e. and define $\phi : t \rightarrow X$ by,

$$\phi(t) = t \cdot u_{tt}. \quad (6)$$

Let f_0 be an eigenfunction with respect to p , with a positive eigenvalue, i.e. $\exists \lambda \in (0, \Lambda < \infty) : p(f_0) = \lambda \cdot f_0$, then

$$u(t) = (1 - \lambda t)^+ f_0, \quad (7)$$

satisfies Eq. (5), where $(q)^+ = \max(0, q)$. This can be verified by taking the time derivative on both sides and using the zero-homogeneity of p . 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 $\phi(t) = \delta(t - \frac{1}{\lambda})f_0$, i.e. ϕ ’s energy is concentrated in a single scale (“frequency”) which corresponds to the eigenvalue of f_0 , $\lambda = \frac{1}{t}$. For a general $f_0 \in X$, this motivates the interpretation of ϕ as a spectral transform of f_0 , where the spectral components are positive eigenfunctions of p , in a similar manner to [5, 7, 13].

We can compute the reconstruction formula, for a general stopping time T , using integration by parts (and assuming $u_t(0)$ is bounded), by

$$\int_0^T \phi(t) dt = tu_t|_0^T - \int_0^T u_t dt = Tu_t(T) - u(T) + f_0 = -Tp(u(T)) - u(T) + f_0.$$

Denoting the residual $R = Tp(u(T)) + u(T)$, the following reconstruction identity holds $f_0 = \int_0^T \phi(t) dt + R$. Similar to gradient flows of one-homogeneous functionals, which are a special case of zero-homogeneous flows, this spectral framework resembles in some sense classical Fourier analysis, e.g. - we can filter f_0 using a transfer function (window) $H : \mathbb{R}^+ \rightarrow \mathbb{R}$ as follows:

$$f_0^{filtered} = \int_0^T H(t)\phi(t) dt + R, \quad (8)$$

where for $H(t) \equiv 1$ (“All-pass filter”) we obtain the reconstruction formula. This is the core capability we use for shape processing in all three methods proposed below. Note that contrary to previous studies, the above properties rely on very mild assumptions regarding the flow. We do not assume the flow minimizes a one-homogeneous functional (and for a finite stopping time, even strict convergence is not necessary, as all diverging components are in R and the reconstruction formula is valid). Our findings are straightforward extensions of observations done in [7], where further discussion takes place, including other key properties, such as orthogonality of the spectral components and decomposition into eigenfunctions by the flow in certain settings.

Modifying Flows for Nonlinear Spectral Processing. Our framework requires a zero-homogeneous flow evolving on a fixed metric (to induce spectral linear decay). 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.

Technical Briefing. A 2D manifold M embedded in 3D is given by its 3 coordinate functions $S_0(\omega_1, \omega_2) = x_0(\omega_1, \omega_2), y_0(\omega_1, \omega_2), z_0(\omega_1, \omega_2)$ where $\omega_1, \omega_2 \in \Omega \subset \mathbb{R}^2$, inducing the intrinsic metric g_0 . M is discretized as a triangular mesh, and S_0 as the vertex coordinates. Gradient and divergence operators $\nabla_{g_0}, \nabla_{g_0} \cdot$ on M are discretized as in [16]. In all modifications of Eq. (2), if $\mathcal{P} < 2$ we replace the magnitude by $\sqrt{|\nabla_g f|^2 + \epsilon^2}$ for stability. Our flows are implemented using semi-implicit time steps, and differ by their operator. The evolving shape at time t is denoted $S(t)$, and $c(t)$ denotes any of $S(t)$ ’s evolving coordinate functions $x(t), y(t), z(t)$.

3.1 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. (3), 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. (2) 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. (5). Each channel evolves separately, hence the name “unpaired coordinates”. We can now perform nonlinear spectral filtering as in Eq. (8), demonstrated in Fig. 2. While processing each of $c = x, y, z$ independently is sub-optimal, this approach has good feature control, as spectral decomposition is applied directly to the embedded shape.

3.2 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, previously suggested by [11], which results in the operator,

$$-p_{M1}(c) := \nabla_{g_0} \cdot \left(\frac{\nabla_{g_0} c}{\sqrt{\sum_{\tilde{c}=x,y,z} |\nabla_{g_0} \tilde{c}|^2}} \right). \quad (9)$$

Note that the metric is fixed as g_0 . We can also verify that the operator is zero-homogeneous. The flow is followed by per-coordinate spectral processing as in Eq. (6), (8) - thus x, y, z inter-correlation is preserved. We obtain good explicit feature control (see Fig. 2).

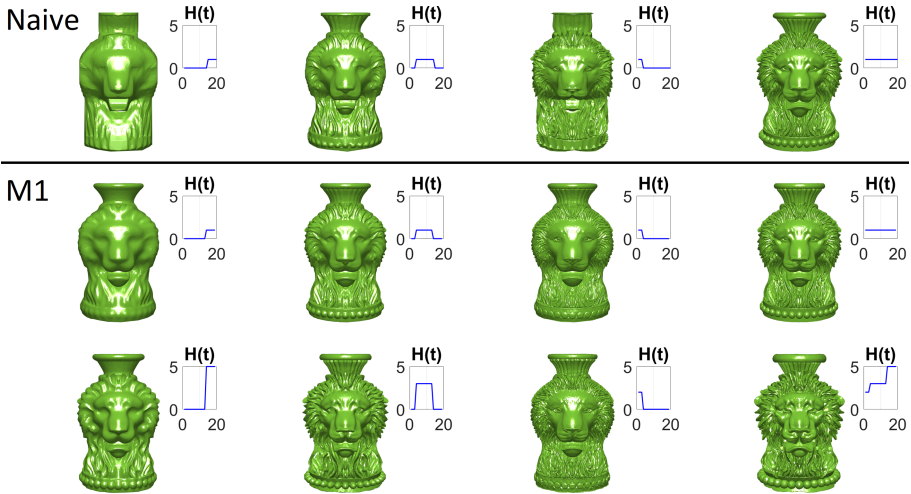


Fig. 2. [Please zoom in] Filtering is applied to the embedded shape, thus magnification and summation of filter-bands magnifies and sums features directly. As t grows, choosing $H(t)$ greater or less than 1, results in amplification or attenuation of finer details. This can be observed in the model’s contour and hair-strands. M1 does not cause the shape to be axis-aligned (squaring effect), as in the naive approach.

3.3 Method 2 (M2): Conformalized \mathcal{P} -Laplace

Here we modify cMCF [17] to our framework, using our conformalized \mathcal{P} -Laplace described below. Our flow inherits cMCF's limb-head smoothing capabilities (Fig. 3), which we then use for shape filtering. The metric of cMCF, $\tilde{g}_t = \sqrt{|g_0^{-1} g_t|} g_0$, is not fixed, and 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 a 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,

$$\tilde{\Delta}_{g,\mathcal{P}}(c) := \sqrt{\frac{|g_0|}{|g_t|}} \nabla_{g_0} \cdot (|\nabla_{g_0} c|^{\mathcal{P}-2} \nabla_{g_0} c). \quad (10)$$

By Eq. (1) we have that $|g_t|$ is absolutely 4-homogeneous, hence $\tilde{\Delta}_{g,\mathcal{P}}$ is $\mathcal{P} - 3$ homogeneous,

$$\tilde{\Delta}_{g,\mathcal{P}}(\alpha c) = \sqrt{\frac{|g_0|}{|\alpha|^4 |g_t|}} \nabla_{g_0} \cdot (|\alpha|^{\mathcal{P}-2} |\nabla_{g_0} c|^{\mathcal{P}-2}) \alpha \nabla_{g_0} c = \frac{\alpha}{|\alpha|^{4-\mathcal{P}}} \tilde{\Delta}_{g,\mathcal{P}}(c).$$

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 [11], yielding the operator,

$$-p_{M2}(c) := \sqrt{\frac{|g_0|}{|g_t|}} \nabla_{g_0} \cdot \left(\sqrt{\sum_{\tilde{c}=x,y,z} |\nabla_{g_0} \tilde{c}|^2} \nabla_{g_0} c \right) \quad (11)$$

The flow is followed by nonlinear spectral filtering, Eq. (8). Editing extremities, a capability inherited from our conformal 3-Laplace flow, is demonstrated in Fig. 4, where extremities are in the form of human limbs and head.

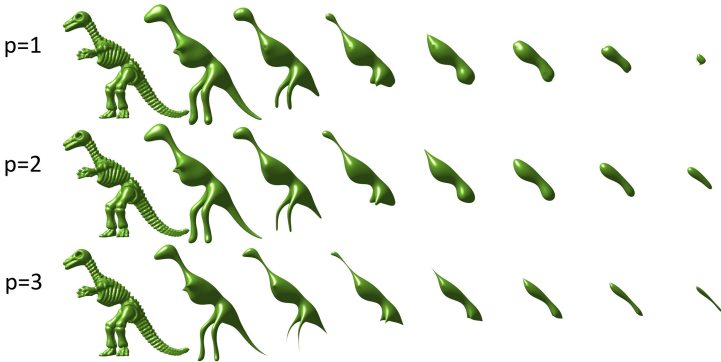


Fig. 3. M2 conformalized \mathcal{P} -Laplace flows. $\mathcal{P} = 2$ is an unscaled version of cMCF. $\mathcal{P} = 1$ is a new conformalized shape TV flow. For $\mathcal{P} = 3$ the flow is zero-homogeneous.

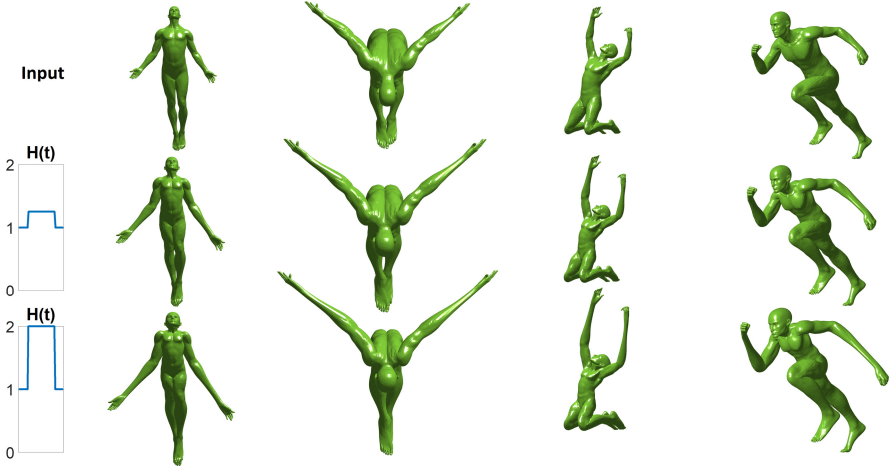


Fig. 4. Shape exaggeration by M2 spectral filtering. The filter inherits its properties from the conformalized 3-Laplace flow (Fig. 3), which translates to interesting limb-head editing. Isometry robustness is demonstrated as well.

3.4 Method 3 (M3): Directional Shape TV

Mesh TV smoothing typically preserve pointy surface points, e.g. tip of chin [12] or ears [11]. Here 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 [18]. 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.

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 makes sure \hat{d} points “outwards”. Finally, the evolving initial surface is represented as $f_0 = \alpha |S - \hat{S}|$. Note that $S = \hat{S} + f_0 \hat{d}$. 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 $\alpha |S - \hat{S}| = (0, 0, f(x, y))$.

We advocate the choice of \hat{S} as a cMCF smoothed version of S , because 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,

$$-p_{M3}(u) := \nabla_{g_0} \cdot \frac{\nabla_{g_0} u}{|\nabla_{g_0} u|}, \quad (12)$$

where $u(t)$ is the evolution of f_0 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_0 is filtered as in Eq. (8), and a filtered shape is obtained by $\hat{S} + f^{filtered} \hat{d}$. Being closely related to spectral TV on images, this method preserves detail well, as demonstrated in Figs. 1, 5, 6. Though inspired by MCF, the proposed flow is substantially different.

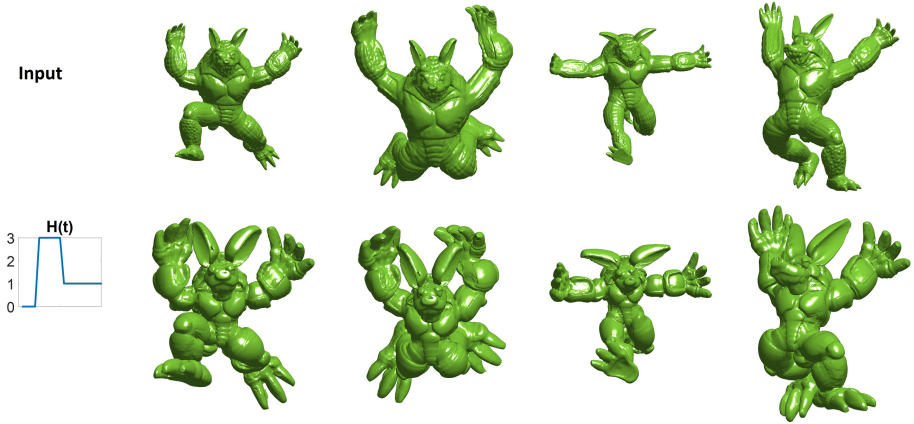


Fig. 5. Shape exaggeration by M3. This filter applies edge preserving smoothing while admitting the requirements for the caricaturization task, as posed by [20].

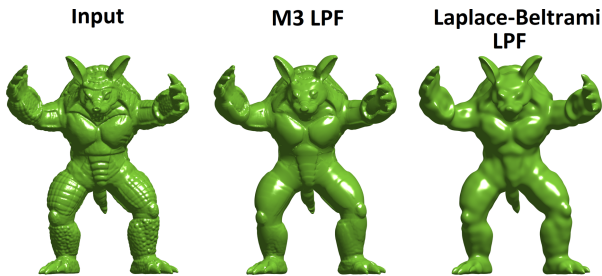


Fig. 6. M3 filtering has distinctive detail-preserving capabilities. All our methods can utilize rough time discretization for runtime efficiency at the expense of reconstruction error - here M3 takes approximately 30 s. The Laplace-Beltrami filtering requires SVD for 2000 eigenvectors of the $[17 \cdot 10^3 \times 17 \cdot 10^3]$ sparse operator matrix, which took us approximately 5 h to obtain (2.5 orders of magnitude slower).

4 Discussion and Conclusion

We propose a general methodology for shape processing based on nonlinear spectral filtering. In this framework we use smoothing flows which satisfy two requirements: zero-homogeneity, and a fixed metric. We process the shape in its embedding space, providing unmediated nonlinear spectral representations and accurate and meaningful filtering capabilities.

To showcase the general concept, three methods are proposed, where spectral processing is done via the same mechanism, Eqs. (6) and (8). M1 is arguably the natural setting for shape spectral TV, and is the most robust and easy to handle. M2 enables the flexible amplification of surface extremities. M3 best entails the well known characteristics of spectral TV on images (such as edge preservation). This is accomplished by restricting the direction of the evolution, generalizing a key component of the TV and Mean Curvature Flow duality. 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. In order to filter by linear eigenfunctions, one first needs to compute the basis for the specific shape. This entails solving SVD of the Laplace-Beltrami operator, which is significantly more costly than our proposed methods (e.g. the computations for Fig. 6 are 30s for M3 vs. 5h for Laplace-Beltrami filtering).

The characteristics of the proposed nonlinear spectral methods are distinct and can carry-over to various applications. Our findings on zero-homogeneous processing are not restricted to shape analysis, and can be implemented even in new neural-network architectures that comply with our requirements. ²

References

1. Andreu, F., Ballester, C., Caselles, V., Mazón, J.M.: Minimizing total variation flow. *Differ. Integral Equ.* **14**(3), 321–360 (2001)
2. Benning, M., et al.: Nonlinear spectral image fusion. In: Lauze, F., Dong, Y., Dahl, A.B. (eds.) *SSVM 2017*. LNCS, vol. 10302, pp. 41–53. Springer, Cham (2017). https://doi.org/10.1007/978-3-319-58771-4_4
3. Bresson, X., Chan, T.F.: Fast dual minimization of the vectorial total variation norm and applications to color image processing. *Inverse Probl. Imag.* **2**(4), 455–484 (2008)
4. Bronstein, A.M., Bronstein, M.M., Kimmel, R.: *Numerical Geometry of Non-Rigid Shapes*. MCS. Springer, New York (2009). <https://doi.org/10.1007/978-0-387-73301-2>
5. Bungert, L., Burger, M., Chambolle, A., Novaga, M.: Nonlinear spectral decompositions by gradient flows of one-homogeneous functionals. *Analysis & PDE* (2019)

² Models: Bust of Queen Nefertiti. gypisches Museum und Papyrussammlung. Model: Trigon art; Stanford armadillo and poses by Belyaev, Yoshizawa, Seidel (2006); Michaels from [4]; various models from LIRIS database.

6. Burger, M., Gilboa, G., Osher, S., Xu, J.: Nonlinear inverse scale space methods. *Commun. Math. Sci.* **4**(1), 179–212 (2006)
7. Burger, M., Gilboa, G., Moeller, M., Eckardt, L., Cremers, D.: Spectral decompositions using one-homogeneous functionals. *SIAM J. Imag. Sci.* **9**(3), 1374–1408 (2016)
8. Burger, M., Osher, S.: A guide to the TV Zoo. In: *Level Set and PDE Based Reconstruction Methods in Imaging*, pp. 1–70. Springer, Cham (2013). https://doi.org/10.1007/978-3-319-01712-9_1
9. Chambolle, A., Caselles, V., Cremers, D., Novaga, M., Pock, T.: An introduction to total variation for image analysis. In: *Theoretical Foundations and Numerical Methods for Sparse Recovery*, vol. 9, no. 263–340, p. 227 (2010)
10. Do Carmo, M.P.: *Differential Geometry of Curves and Surfaces: Revised and Updated Second Edition*. Courier Dover Publications, Mineola (2016)
11. Elmoataz, A., Lezoray, O., Boughleux, S.: Nonlocal discrete regularization on weighted graphs: a framework for image and manifold processing. *IEEE Trans. Image Process.* **17**(7), 1047–1060 (2008)
12. Fumero, M., Möller, M., Rodolà, E.: Nonlinear spectral geometry processing via the TV transform. *ACM Trans. Graph. (TOG)* **39**(6), 1–16 (2020)
13. Gilboa, G.: A total variation spectral framework for scale and texture analysis. *SIAM J. Imag. Sci.* **7**(4), 1937–1961 (2014)
14. Gilboa, G.: A spectral approach to total variation. In: Kuijper, A., Bredies, K., Pock, T., Bischof, H. (eds.) *SSVM 2013*. LNCS, vol. 7893, pp. 36–47. Springer, Heidelberg (2013). https://doi.org/10.1007/978-3-642-38267-3_4
15. Hait, E., Gilboa, G.: Spectral total-variation local scale signatures for image manipulation and fusion. *IEEE Trans. Image Process.* **28**(2), 880–895 (2019)
16. Jacobson, A., Panozzo, D.: *Libigl: prototyping geometry processing research in C++* (2017). <https://libigl.github.io/tutorial/#chapter-2-discrete-geometric-quantities-and-operators>
17. Kazhdan, M., Solomon, J., Ben-Chen, M.: Can mean-curvature flow be modified to be non-singular? In: *Computer Graphics Forum*, vol. 31, pp. 1745–1754. Wiley, Hoboken (2012)
18. Kimmel, R., Malladi, R., Sochen, N.: Images as embedded maps and minimal surfaces: movies, color, texture, and volumetric medical images. *Int. J. Comput. Vis.* **39**(2), 111–129 (2000)
19. Mantegazza, C.: *Lecture Notes on Mean Curvature Flow*, vol. 290. Springer, Basel (2011). <https://doi.org/10.1007/978-3-0348-0145-4>
20. Sela, M., Aflalo, Y., Kimmel, R.: Computational caricaturization of surfaces. *Comput. Vis. Image Underst.* **141**, 1–17 (2015)
21. Sorkine, O., Cohen-Or, D., Lipman, Y., Alexa, M., Rössl, C., Seidel, H.P.: Laplacian surface editing. In: *Proceedings of the 2004 Eurographics/ACM SIGGRAPH Symposium on Geometry Processing*, pp. 175–184 (2004)
22. Taubin, G.: A signal processing approach to fair surface design. In: *Proceedings of 22nd Annual Conference on Computer Graphics and Techniques*, pp. 351–358 (1995)
23. Vallet, B., Lévy, B.: Spectral geometry processing with manifold harmonics. In: *Computer Graphics Forum*, vol. 27, pp. 251–260. Wiley, Hoboken (2008)
24. Wetzler, A., Aflalo, Y., Dubrovina, A., Kimmel, R.: The Laplace-Beltrami operator: a ubiquitous tool for image and shape processing. In: Hendriks, C.L.L., Borgefors, G., Strand, R. (eds.) *ISMM 2013*. LNCS, vol. 7883, pp. 302–316. Springer, Heidelberg (2013). https://doi.org/10.1007/978-3-642-38294-9_26