

Natural Boundary Conditions for Smoothing in Geometry Processing

ODED STEIN, Columbia University, USA

EITAN GRINSPUN, Columbia University, USA

MAX WARDETZKY, Universität Göttingen, Germany

ALEC JACOBSON, ETH Zurich, Switzerland, Columbia University, USA, and University of Toronto, Canada

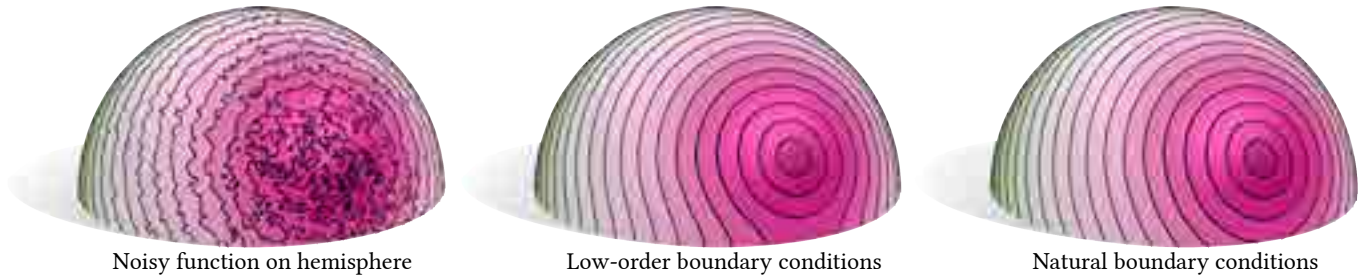


Fig. 1. Smoothing a noisy function with the Laplacian energy and common low-order boundary conditions (zero Neumann) introduces a bias at the boundary: isolines exit perpendicularly. We propose using a different smoothness energy, the Hessian energy, whose natural boundary conditions avoid this bias.

In geometry processing, smoothness energies are commonly used to model scattered data interpolation, dense data denoising, and regularization during shape optimization. The squared Laplacian energy is a popular choice of energy and has a corresponding standard implementation: squaring the discrete Laplacian matrix. For compact domains, when values along the boundary are not known in advance, this construction *bakes in* low-order boundary conditions. This causes the geometric shape of the boundary to strongly bias the solution. For many applications, this is undesirable. Instead, we propose using the squared Frobenius norm of the Hessian as a smoothness energy. Unlike the squared Laplacian energy, this energy's *natural boundary conditions* (those that best minimize the energy) correspond to meaningful high-order boundary conditions. These boundary conditions model free boundaries where the shape of the boundary should not bias the solution locally. Our analysis begins in the smooth setting and concludes with discretizations using finite-differences on 2D grids or mixed finite elements for triangle meshes. We demonstrate the core behavior of the squared Hessian as a smoothness energy for various tasks.

CCS Concepts: • **Mathematics of computing** → **Discretization; Partial differential equations; Numerical differentiation**; • **Computing methodologies** → **Animation; Mesh geometry models**;

Additional Key Words and Phrases: Geometry processing, smoothing, denoising, Hessian energy, Laplacian energy

Authors' addresses: Oded Stein, Columbia University, APAM, 500 W 120th St, Room 200, New York, NY, 10027, USA, oded.stein@columbia.edu; Eitan Grinspun, Columbia University, Computer Science Department, 500 W 120th St, Room 450, New York, NY, 10027, USA, eitan@cs.columbia.edu; Max Wardetzky, Universität Göttingen, Institute of Num. and Appl. Math, Lotzestr. 16-18, Göttingen, 37083, Germany, wardetzky@math.uni-goettingen.de; Alec Jacobson, ETH Zurich, Switzerland, Columbia University, USA, University of Toronto, Bahen Centre for Information Technology, 40 St. George Street, Room 5266, Toronto, ON, M5S 2E4, Canada, alecjacobson@gmail.com.

© 2010 Copyright held by the owner/author(s).
This is the author's version of the work. It is posted here for your personal use. Not for redistribution. The definitive Version of Record was published in *ACM Transactions on Graphics*, https://doi.org/0000001.0000001_2

ACM Reference Format:

Oded Stein, Eitan Grinspun, Max Wardetzky, and Alec Jacobson. 2010. Natural Boundary Conditions for Smoothing in Geometry Processing. *ACM Trans. Graph.* 9, 4, Article 39 (March 2010), 13 pages. https://doi.org/0000001.0000001_2

1 INTRODUCTION

Smoothness energies are commonly used in graphics, geometry processing and image processing to model deformations, denoise densely sampled measurements, and interpolate sparse scattered data. Compared to those involving first derivatives, smoothness energies involving second derivatives are by definition more expressive, but also more challenging to control. For example, minimizing the squared-norm of the gradient of an unknown function u integrated over a bounded domain $\Omega \subset \mathbb{R}^2$

$$E_{\nabla^2}(u) = \frac{1}{2} \int_{\Omega} \|\nabla u\|^2 dA, \quad (1)$$

results in a second-order Laplace equation with *only* one set of prescribable boundary conditions (either fixed values, normal derivatives, or linear combinations thereof).

Meanwhile, minimizing the *squared Laplacian energy*

$$E_{\Delta^2}(u) = \frac{1}{2} \int_{\Omega} (\Delta u)^2 dA, \quad (2)$$

results in a fourth-order *bi*-Laplace equation with many different combinations of prescribable boundary conditions (values and normal derivatives; values and Laplacians; normal derivatives and Laplacians; etc.). This greater expressive power comes with greater responsibility during modeling.

This greater responsibility manifests itself in the need for carefully choosing *boundary conditions* for higher order energies, such as the squared Laplacian energy. Indeed, lacking a careful and systematic treatment of boundary conditions can lead to undesirable

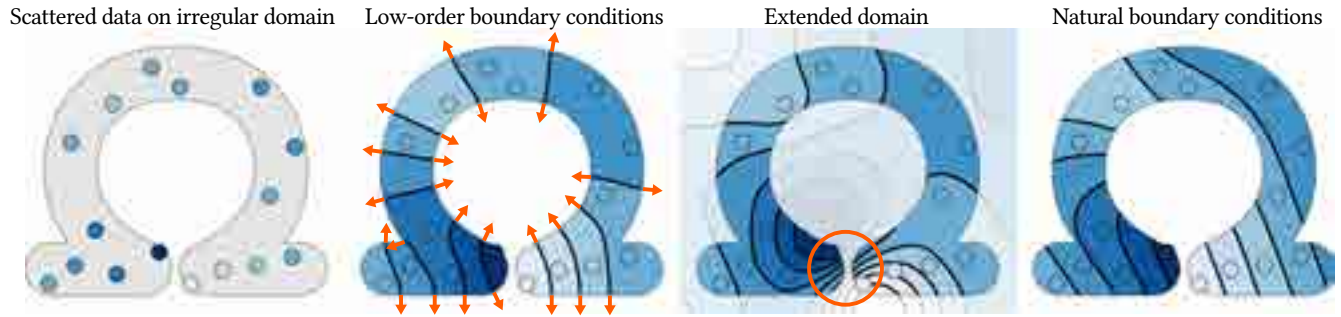


Fig. 2. When interpolating scattered data over a bounded domain, Neumann conditions introduce a noticeable bias at the boundary (orange arrows). Meanwhile, ignoring the boundary by interpolating over ambient space allows *bleeding* of data onto geodesically distant regions (orange circle).

and unintuitive artifacts in geometry processing applications. We here refer to boundary conditions that only depend on the function value and its first derivatives, such as Dirichlet and Neumann boundary conditions, as low-order. Boundary conditions involving higher derivatives are referred to as high-order. Consider for example the problem of data smoothing. Enforcing Neumann boundary conditions for the squared Laplacian energy leads to noticeable bias along the boundary which is often unintentional, as illustrated in Fig. 1 (middle). Likewise, consider the problem of scattered data interpolation, such as in Fig. 2. Here Neumann boundary conditions will cause the interpolation to be strongly influenced by the boundary’s shape and isolines will be perpendicular to the boundary (Fig. 2, orange arrows). One option would be to simply ignore the given domain’s boundary and instead solve over a naive extension into the surrounding space. As a consequence, however, data will *bleed* across the original domain boundaries (Fig. 2, orange circle).

In this paper, we offer a study of the effect of boundary conditions for second-order energies derived from the biharmonic equation. In particular, we show limitations of the popular squared Laplacian energy, and we propose an alternative second-order energy: the squared Frobenius norm of the Hessian or simply *squared Hessian energy*:

$$E_{H^2}(u) = \frac{1}{2} \int_{\Omega} \|\mathbf{H}_u\|_F^2 dA, \quad (3)$$

where $\mathbf{H}_u \in \mathbb{R}^{2 \times 2}$ is the symmetric matrix of second-order partial derivatives. While this energy appears frequently in the mechanics and image processing literatures, it has not been applied and analyzed as a smoothness energy for geometry processing on irregular domains.

If both Dirichlet and Neumann boundary conditions are enforced simultaneously, then minimizers of E_{H^2} and E_{Δ^2} will be identical. However, these two energies offer dramatically different behavior for “free” or *natural* boundaries where no explicit boundary conditions are enforced. *Natural boundary conditions* are those that, among all possible boundary conditions, minimize the given energy.

Unlike the natural boundary conditions of the squared Laplacian energy, those of the squared Hessian energy admit well-behaved solutions without strong bias near the boundary (see Fig. 3). We

show the implications of this when modeling smoothness in geometry processing. We derive these boundary conditions and analyze their effects in detail. The advantages of the squared Hessian energy are especially pronounced for irregular domains (e.g., non-convex domains with arbitrarily shaped, non-axis-aligned boundaries).

To demonstrate practical results for typical applications of smoothness energies, we discretize the squared Hessian energy with natural boundary conditions following the same steps that have been used to discretize the squared Laplacian energy previously, using finite differences on 2D grids and mixed finite-elements on triangle meshes.

Alongside our core contribution of providing a boundary-insensitive smoothness energy for geometry processing tasks on irregular domains, we also present a novel L_1 -norm minimization for piecewise planar data reconstruction as well as a smooth understanding of previous discrete methods for linear subspace design for real-time shape deformation.

2 BACKGROUND

Before discussing related works specifically, we establish concretely what we mean by *natural boundary conditions*.

2.1 Definition: natural boundary conditions

In the calculus of variations, the term natural boundary conditions refers specifically to boundary conditions appearing in the Euler-Lagrange equation of an energy minimization problem that were not imposed beforehand (see [Gelfand and Fomin 1963, p. 26] or [Giaquinta and Hildebrandt 1996, p. 34]). This definition is general in that it applies to any energy or partially constrained energy. Just like other types of boundary conditions, it is possible to witness natural boundary conditions on one part of a domain’s boundary and explicitly enforce different boundary conditions on another part.

Natural boundary conditions achieve the lowest energy among all possible boundary conditions for a given minimization problem. Consequently, natural boundary conditions heavily depend on the energy being minimized. As we will see, the natural boundary conditions of the squared gradient energy (Dirichlet energy) in Equation (1), the squared Laplacian energy in Equation (2), and the Equation (3) are all strikingly different.

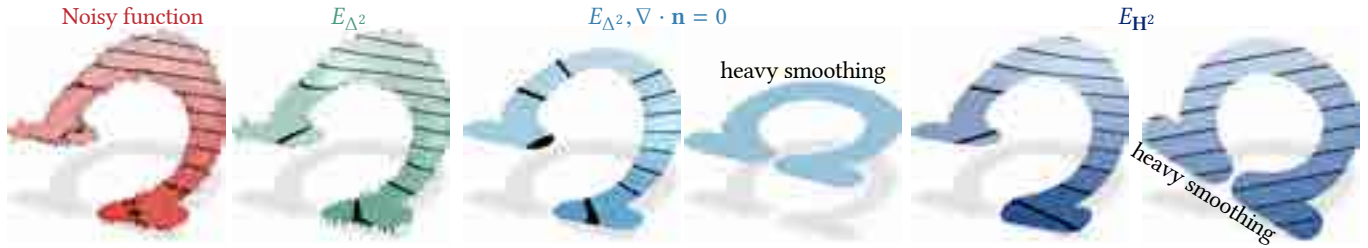


Fig. 3. Smoothing a function over an omega-shaped domain. The high-order natural boundary conditions of E_{Δ^2} maintain noisy boundaries (green). Previous works use low-order boundary conditions, but this biases the solution to *follow* the boundary even if the input data does not (light blue). Smoothing converges to a constant function. In contrast, the high-order natural boundary conditions of E_{H^2} produce smooth boundary values without bias and converges to a *linear* function (dark blue).

2.2 Related Work

The natural boundary conditions of E_{Δ^2} and E_{H^2} were considered by Courant and Hilbert [1924]. Since then, second-derivative energies have been considered in various fields, including data, image, and geometry processing. Earlier work also relates to ours in how the energy, boundary conditions, and domain are chosen, and in their application.

2.2.1 Geometry Processing. High-order smoothness energies are a popular and powerful tool. Discretizations of the squared Laplacian energy have been used for surface fairing [Desbrun et al. 1999], smooth surface displacements [Andrews et al. 2011; Botsch and Kobbelt 2004; Sorkine et al. 2004; Zhou et al. 2005], smooth geodesic distance computation [Lipman et al. 2010], data smoothing [Weinkauff et al. 2010], and regularization during other high-level operations [Cao et al. 2015, 2014; Huang et al. 2006; Jones et al. 2016; Zhou et al. 2010]. These and other works *square* the discrete cotangent Laplacian for triangle meshes [Pinkall and Polthier 1993]. This Laplacian matrix L is *constructed* assuming zero Neumann boundary conditions (i.e., for a boundary edge ij with opposite angle α_{ij} , $L_{ij} = \cot \alpha_{ij}$). Squaring this matrix effectively *bakes* in these conditions. Jacobson et al. [2010] confirm that this construction agrees with a convergent mixed finite element discretization of minimizing E_{Δ^2} *subject to* zero Neumann boundary conditions ($\nabla u \cdot \mathbf{n} = 0$).

Squaring L is an intuitive and easy way to handle unconstrained boundaries when employing E_{Δ^2} . However, we show that zero Neumann boundary conditions are *not natural* for E_{Δ^2} according to the mathematical definition in Section 2.1. Perhaps surprisingly, the true natural boundary conditions of E_{Δ^2} are *not* useful in general (see Fig. 3). Meanwhile, the “not natural” zero Neumann conditions are clearly quite useful, as evidenced by the sheer number of geometry processing methods that impose them and achieve great results for their respective applications [Finch et al. 2011; Jacobson et al. 2011, 2012; Kavan et al. 2011; Landreneau and Schaefer 2010; Lipman et al. 2010; Rustamov 2011; Sýkora et al. 2014; Weber et al. 2012; Weinkauff et al. 2010]. And yet, we will present evidence that the natural boundary conditions of E_{H^2} can be more useful than either the zero Neumann or natural boundary conditions of E_{Δ^2} in contexts where boundaries should minimally influence the solution.

Natural boundary conditions for energies involving *first* derivatives are found frequently in geometry processing or related fields

such as computer animation [Batty 2010; Bojsen-Hansen and Wojtan 2016]. When parameterizing surfaces, natural boundary conditions for the as-conformal-as-possible energy remove the burden of prescribing values for cut-boundaries [Cohen-Steiner and Desbrun 2002; Desbrun et al. 2002; Lévy et al. 2002; Mullen et al. 2008; Springborn et al. 2008]. The boundary conditions imposed by Desbrun et al. [2002] have been applied to *different* energies to which they are no longer natural per Section 2.1, though they nonetheless serve the application [Gingold et al. 2006].

Previous works in geometry processing have considered how to discretize the Hessian of a scalar function defined on a triangulated surface [de Goes et al. 2014; Tosun et al. 2007]. These methods do not explore minimization of E_{H^2} or its natural boundary conditions.

Wang et al. [2015] square a modified cotangent Laplacian to design linear subspaces for cartoon animation; they do not explicitly discuss boundary conditions. We will consider the *discrete* energy in [Wang et al. 2015] and its relationship to others built using discrete exterior calculus [Fisher et al. 2007] and non-conforming edge-based elements [Bergou et al. 2006] in Section 4.5. We show that — in contrast to these previous discrete energies — the null space of E_{H^2} *contains and only contains* affine functions, the necessary and sufficient condition for linear subspace design as set forth by Wang et al.

In the infinite Euclidean domain, biharmonic radial basis functions (a.k.a, thin-plate splines) minimize E_{H^2} and form a useful basis for various scattered data interpolation problems [Bookstein 1989]. Botsch and Kobbelt [2005] applied these globally supported functions to shape deformation and modelling by densely sampling selected surface regions. Free surface boundaries are not *noticed* any differently from the surface interior as both are embedded in a deformation of the entire Euclidean space (see Fig. 4).

2.2.2 Energies and boundary conditions in other fields.

Second-derivative smoothness energies such as E_{Δ^2} and E_{H^2} are found frequently in denoising, restoration, inpainting, image enhancement, and domain reduction applications. Works in image processing that specifically consider E_{H^2} [Didas and Weickert 2004; Lefkimmatis et al. 2012; Lysaker and Tai 2006; Roth 2009; Terzopoulos 1984] or a broader class of energies including E_{H^2} as a special case [Didas et al. 2009; Lefkimmatis et al. 2012] impose low-order boundary conditions, with some notable exceptions: Terzopoulos [1984; 1988] was the first to employ the natural boundary conditions of

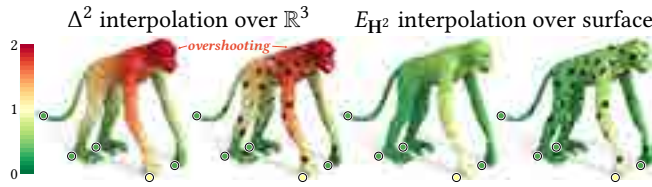


Fig. 4. Scattered data interpolation using biharmonic radial basis functions in \mathbb{R}^3 ignores the embedded surface and therefore also surface boundaries. In contrast, interpolation by minimizing E_{H^2} on the surface is “shape aware”, yet not qualitatively disturbed by open boundaries.

E_{H^2} in the vision and graphics literature, in the context of surface reconstruction from images. Didas et al. [2004; 2009] invoked the natural conditions for image denoising, whereas Lefkimiatis et al. [2012] removed boundaries altogether via periodic or reflexive identification. When smoothing over the convex image rectangle, free boundaries have diminished influence on the behavior of the solution. Applications such as inpainting frequently consider non-convex, irregular image subregions, but enforce low-order boundary conditions explicitly to ensure value and derivative continuity [Georgiev 2004].

Other works consider related energies, such as the L_1 -norm of the Hessian [Lysaker et al. 2003; Steidl 2006; Yuan et al. 2009] (a generalization of *total variation* [Rudin et al. 1992]), the energy E_{Δ^2} [Steidl et al. 2005] or variants thereof [You and Kaveh 2000]. We show how the L_1 -norm can analogously apply to smoothing geometric data.

Other authors formulate fourth-order equations directly, not via an energy of second-derivatives. For instance, Liu et al. [2015] explicitly enforce second- and third-order boundary conditions ($\nabla u \cdot \mathbf{n} = \nabla \Delta u \cdot \mathbf{n} = 0$) on a fourth-order equation similar to the bi-Laplace equation. As we will see in Section 4, this set of boundary conditions is not *natural* to either E_{Δ^2} or E_{H^2} , however they are frequently imposed on E_{Δ^2} for geometry processing problems on surfaces with boundaries.

E_{Δ^2} and E_{H^2} have been used for the spectral embedding of high dimensional data [Belkin and Niyogi 2004; Donoho and Grimes 2003]. A generalization of E_{H^2} was used by Steinke and Hein [2009] for regression, using the energy’s natural boundary conditions (therein referred to as implicit).

In physical systems, natural boundary conditions typically have an intuitive explanation: they explain the behavior of free boundaries. In elastic beam theory, natural boundary conditions are higher-order and appear for biharmonic-like equations (see, e.g., [Li 2000]). Similar boundary conditions also appear in Kirchhoff plate bending theory [Felippa 2017]. We now leverage this physical intuition in one-dimension before discussing smoothing in two-dimensions and on surfaces.

3 ONE DIMENSION

As a didactic exercise, let us derive natural boundary conditions for the one-dimensional squared second derivative energy integrated

over the unit line segment:

$$\min_u \frac{1}{2} \int_0^1 (u'')^2 dx, \quad (4)$$

where $u' = du/dx$, $u'' = d^2u/dx^2$ and so on. Throughout, we appeal to the physical metaphor of a bending bar (see Fig. 5).

In practice, we minimize this energy subject to various boundary conditions or additional energy terms. However, applying the calculus of variations to the energy’s raw form will better illustrate natural boundary conditions. In the absence of explicit boundary conditions, then u is a minimizer of this energy if *any* infinitesimal variation εv does not change its energy:

$$0 = \left. \frac{d}{d\varepsilon} \right|_{\varepsilon=0} \frac{1}{2} \int_0^1 ((u + \varepsilon v)'')^2 dx = \int_0^1 u'' v'' dx \quad (5)$$

for all sufficiently smooth functions v . In particular, v is *not* constrained at the boundary. Integration by parts yields

$$0 = \int_0^1 u'''' v dx - [u'''' v]_0^1 + [u'' v']_0^1 \quad (6)$$

for all v . By judiciously testing with v functions, we conclude that $u'''' = 0$ on the interior interval $(0, 1)$ and $u'' = u'''' = 0$ at the boundary $\{0, 1\}$. We began without explicit boundary conditions, so these are the *natural boundary conditions*.

We observe that functions in the energy’s null space satisfy that energy’s natural boundary conditions – any linear function u measures zero energy. Correspondingly, we can interpret the natural boundary conditions as “the solution should be linear at the boundary” ($u'' = 0$).

We can repeat this process even if we fix *some but not all* boundary conditions explicitly. For example, if we fix $u(0) = 0$ then the remaining natural boundary conditions will be $u''(0) = 0$ and $u''(1) = u''''(1) = 0$ leaving any linear function passing through origin as a solution. If we also fix $u'(0) = 0$ then the boundary at $x = 0$ is fully constrained and natural boundary conditions only appear at $x = 1$, namely $u''(1) = u''''(1) = 0$. In this case, the problem has a unique albeit trivial solution: $u = 0$, the only linear function passing through the origin with zero slope.

From a mechanics perspective, natural boundary conditions correspond to zero forces and zero moments at the boundary. Fig. 5 shows an example of a uniform load applied to a bending bar (i.e., $\min_u \int_0^1 (u'')^2 + u dx$) with various explicit boundary conditions and each problem’s resulting set of natural boundary conditions.

In two dimensions, there will be many different ways to construct a smoothness energy out of squared second derivatives. Different energies might even behave the same when boundaries are sufficiently fixed, but imply *different* natural boundary conditions in the presence of free boundaries.

4 TWO DIMENSIONS

Given a flat 2D domain $\Omega \subset \mathbb{R}^2$ with a possibly intricate yet smooth boundary, we consider two intimately related energies that measure second-order variations of a function $u : \Omega \rightarrow \mathbb{R}$: the *squared Laplacian energy* $E_{\Delta^2}(u)$ of Equation (2) and the *squared Hessian energy* $E_{H^2}(u)$ of Equation (3).

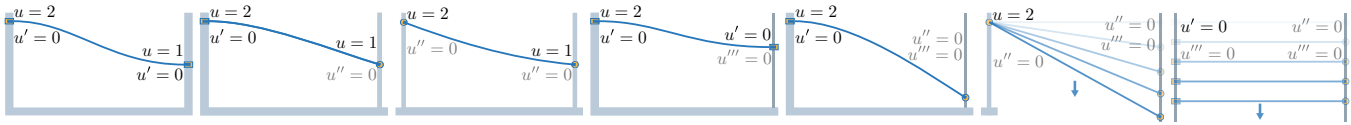


Fig. 5. A bending bar experiences a uniform load subject to various explicit boundary conditions (black, e.g., $u = 2$); in their absence, *natural* boundary conditions emerge (gray, e.g., $u'' = 0$). This physical problem has a unique solution when at least two low-order boundary conditions are imposed; otherwise, the problem becomes underconstrained and the solution tends toward infinite values. Calculus of variation shows that a uniform load does not change the natural boundary conditions.

4.1 Green's Identities

We will employ various *Green's identities* for functions in \mathbb{R}^2 to move derivatives across inner products by introducing a boundary term. Starting with the classic identity:

$$\int_{\Omega} (\nabla u \cdot \nabla v + u \Delta v) \, dA = \oint_{\partial\Omega} u \nabla v \cdot \mathbf{n} \, ds, \quad (7)$$

Replacing u in Eq. (7) with Δu :

$$\int_{\Omega} (\nabla \Delta u \cdot \nabla v + \Delta u \Delta v) \, dA = \oint_{\partial\Omega} \Delta u \nabla v \cdot \mathbf{n} \, ds, \quad (8)$$

Replacing v in Eq. (7) with Δv :

$$\int_{\Omega} (\nabla u \cdot \nabla \Delta v + u \Delta^2 v) \, dA = \oint_{\partial\Omega} u \nabla \Delta v \cdot \mathbf{n} \, ds, \quad (9)$$

Replacing u and v in Eq. (7) with ∇u^T and ∇v^T :

$$\int_{\Omega} (\mathbf{H}_u : \mathbf{H}_v + \nabla u \cdot \nabla \Delta v) \, dA = \oint_{\partial\Omega} \nabla u^T \mathbf{H}_v \mathbf{n} \, ds, \quad (10)$$

Replacing u and ∇v in Eq. (7) with a vector \mathbf{u} and matrix \mathbf{V} :

$$\int_{\Omega} (\nabla \mathbf{u} : \mathbf{V} + \mathbf{u} \cdot (\nabla \cdot \mathbf{V})) \, dA = \oint_{\partial\Omega} \mathbf{u}^T \mathbf{V} \mathbf{n} \, ds, \quad (11)$$

where $\mathbf{X} : \mathbf{Y} := \text{tr}(\mathbf{X}^T \mathbf{Y})$ computes the Frobenius inner product which generates the Frobenius norm $\mathbf{X} : \mathbf{X} = \|\mathbf{X}\|_F^2$, and $\nabla \cdot \mathbf{X}$ computes the matrix divergence of \mathbf{X} .

4.2 Equivalence up to boundary conditions

We first show that – regardless of boundary conditions – minimizers of both E_{Δ^2} and $E_{\mathbf{H}^2}$ satisfy the biharmonic equation ($\Delta^2 u = 0$) in the interior of the domain. To avoid redundant derivations, we introduce a parameter α (cf. [Courant and Hilbert 1924; Terzopoulos 1984]) so that setting $\alpha = 1$ gives $E_{\mathbf{H}^2}$ and $\alpha = 0$ gives E_{Δ^2} :

$$\frac{1}{2} \int_{\Omega} \left((1 - \alpha) (\Delta u)^2 + \alpha \|\mathbf{H}_u\|_F^2 \right) \, dA. \quad (12)$$

As in Section 3, u is a minimizer if adding *any* infinitesimal variation εv will not change its energy. Thus,

$$0 = \int_{\Omega} \left((1 - \alpha) \Delta u \Delta v + \alpha \mathbf{H}_u : \mathbf{H}_v \right) \, dA \quad (13)$$

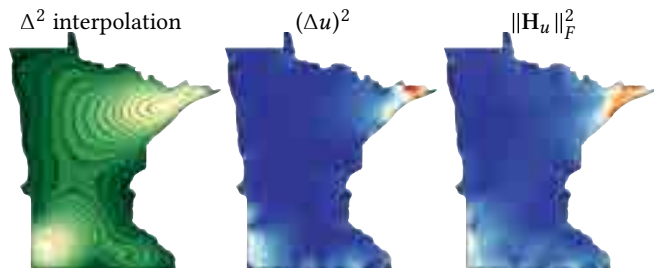


Fig. 6. Subject to fixed boundary values and normal derivatives, the same biharmonic interpolation minimizes *both* E_{Δ^2} and $E_{\mathbf{H}^2}$, despite each energy measuring a different local quantities.

for all v . Applying Equations (8) and (10) this becomes

$$0 = \int_{\Omega} \left(-(1 - \alpha) \nabla \Delta u \cdot \nabla v - \alpha \nabla v \cdot \nabla \Delta u \right) \, dA + \oint_{\partial\Omega} \left((1 - \alpha) \Delta u \nabla v \cdot \mathbf{n} + \alpha \nabla v^T \mathbf{H}_u \mathbf{n} \right) \, ds, \quad (14)$$

for all v . Then applying Equation (9), we arrive at:

$$0 = \int_{\Omega} v \Delta^2 u \, dA + \oint_{\partial\Omega} \left(-v \nabla \Delta u \cdot \mathbf{n} + (1 - \alpha) \Delta u \nabla v \cdot \mathbf{n} + \alpha \nabla v^T \mathbf{H}_u \mathbf{n} \right) \, ds, \quad (15)$$

for all v . Because this equality must hold for any choice of v , u must be *biharmonic* in the interior, regardless of α :

$$\Delta^2 u = 0 \text{ on } \Omega. \quad (16)$$

Biharmonic functions are uniquely determined by boundary values and normal derivatives [Evans 1998, pp. 345], therefore *minimizers* of E_{Δ^2} and $E_{\mathbf{H}^2}$ on \mathbb{R}^2 will be identical when *explicit low-order boundary conditions* are prescribed, i.e., fixing Dirichlet ($u = f$) as well as Neumann ($\nabla u \cdot \mathbf{n} = g$) conditions. Indeed, enforcing such boundary conditions during calculus of variations requires considering variations such that $v = \nabla v \cdot \mathbf{n} = 0$ on the boundary. This immediately implies that $\nabla v = 0$ on the boundary and correspondingly that the entire boundary integral in Equation (15) vanishes. In conclusion, when modeling smoothness for boundary-value interpolation problems, these energies are interchangeable despite measuring different local quantities (see Fig. 6).

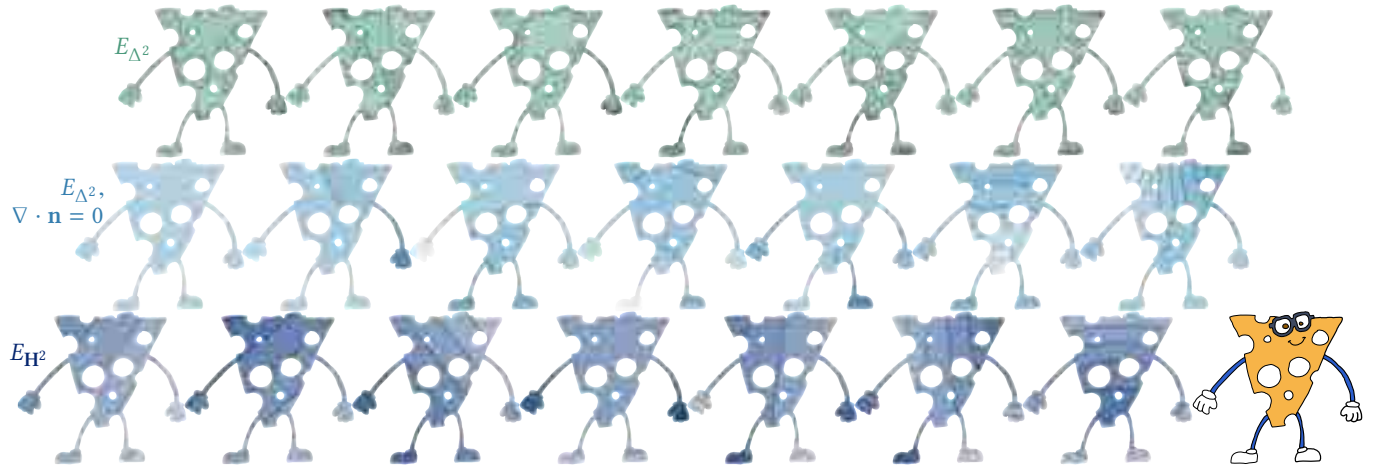


Fig. 7. Ten lowest frequency modes of E_{Δ^2} with natural boundary conditions on a 2D *Cheeseman* include random harmonic functions with high oscillation near the boundary (top). Adding zero Neumann boundary conditions to E_{Δ^2} causes modes to reproduce familiar Laplacian eigenfunctions (middle). In contrast, modes of E_{H^2} with natural boundary conditions include affine functions and other low frequency functions without constrained normal derivatives (bottom).

4.3 Natural boundary conditions

In the absence of explicit boundary conditions, the variations v in Equation (15) include all sufficiently smooth functions. We choose from them specific variations to expose the natural boundary conditions of the parameterized family of energies in Equation (12). First, consider all v that vanish on the boundary but whose gradient exists only in the normal direction ($v = 0$ and $\nabla v = g\mathbf{n}|_{\partial\Omega}$, for some arbitrary function g). Second-order natural boundary conditions must hold:

$$(1 - \alpha)\Delta u + \alpha \mathbf{n}^T \mathbf{H}_u \mathbf{n} = 0 \text{ on } \partial\Omega. \quad (17)$$

Additionally, considering all v with zero normal derivative without restricting the value of v along the boundary ($\nabla v \cdot \mathbf{n} = 0|_{\partial\Omega}$), we witness third-order boundary conditions must hold:

$$\nabla \Delta u \cdot \mathbf{n} + \alpha \nabla \left(\mathbf{t}^T \mathbf{H}_u \mathbf{n} \right) \cdot \mathbf{t} = 0 \text{ on } \partial\Omega. \quad (18)$$

Both conditions depend on α , and therefore the natural boundary conditions for the squared Laplacian and squared Hessian energies indeed differ (ref. [Courant and Hilbert 1924]). We now build an intuition for these conditions for canonical choices of α and contrast their behavior when modeling smoothness on domains with free boundaries.

4.4 Contrasting natural boundary behavior

For the squared Laplacian energy E_{Δ^2} ($\alpha = 0$), unconstrained minimizers of E_{Δ^2} ($\alpha = 0$) will be biharmonic on the interior ($\Delta^2 u = 0$) and satisfy natural boundary conditions forcing the solution to be harmonic along the boundary ($\Delta u = 0$) and continue to be harmonic across the boundary ($\nabla \Delta u \cdot \mathbf{n} = 0$). Consistent with our understanding from Section 2.1, these boundary conditions agree with the energy's null space: any harmonic function u measures zero energy $E_{\Delta^2}(u) = 0$.

For the squared Hessian energy E_{H^2} ($\alpha = 1$), unconstrained minimizers are also biharmonic on the interior ($\Delta^2 u = 0$), but satisfy

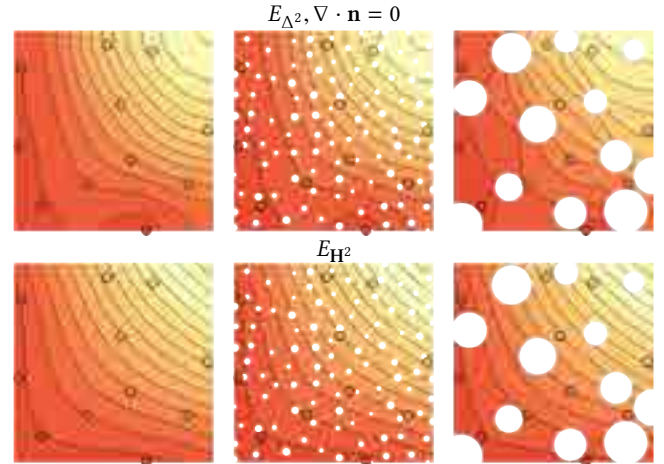


Fig. 8. Interpolating scattered samples of a quadratic function over a square domain without holes (left) and with holes (center and right). When minimizing E_{Δ^2} with Neumann boundary conditions, the presence and shape of boundaries effects the solution. With E_{H^2} , boundaries are less noticeable.

different natural boundary conditions forcing the function to be linear in the normal direction ($\mathbf{n}^T \mathbf{H}_u \mathbf{n} = 0$) and after that continue with low variation along the boundary $\nabla \Delta u \cdot \mathbf{n} + \nabla \left(\mathbf{t}^T \mathbf{H}_u \mathbf{n} \right) \cdot \mathbf{t} = 0$. Again, this agrees with the null space: all linear functions measure zero energy.

The null space of E_{Δ^2} is infinite dimensional: for any boundary values, there exists a harmonic function interpolating them. This renders the E_{Δ^2} with natural boundary conditions rather useless as a smoothness energy. For data smoothing, noisy data along the boundary will simply remain, since there exists an interpolating harmonic function that can be added without affecting the energy (see Fig. 3, *left*).

In contrast, the linear functions in the null space of E_{H^2} are finite dimensional and intuitive. For data smoothing, this null space models “in the absence of any other information, fit a linear function” (see Fig. 3, *right*).

Low-frequency eigenfunctions of smoothness energies are a standard way to construct a smooth low-frequency function space in geometry processing. Typically, previous works use the low-order Dirichlet energy in Equation (1) with its natural boundary conditions ($\nabla u \cdot \mathbf{n} = 0$) [Hildebrandt et al. 2011; Zhang et al. 2007]. Such modal analysis on the higher-order E_{Δ^2} with its natural boundary conditions is not practical because low-frequency modes are polluted with arbitrary harmonic functions (see Fig. 7, *top*). Explicitly enforcing zero Neumann boundary conditions $\nabla u \cdot \mathbf{n} = 0$ on E_{Δ^2} alleviates this but simply results in the same spectrum as E_{∇^2} (up to squared eigenvalues, [Lipman et al. 2010]). In contrast, the low frequency modes of E_{H^2} provide a new, smooth, and well-behaved basis (see Fig. 7).

4.4.1 Mixing with explicit boundary conditions. We examine how explicitly enforcing certain boundary conditions changes the remaining natural boundary conditions. For example, to explicitly enforce values of u along the boundary (Dirichlet conditions), we must assume that our test function v vanishes on the boundary ($v = 0|_{\partial\Omega}$ in Equations 13-15. Only second-order natural boundary conditions remain. For E_{Δ^2} , the additional natural boundary conditions are $\Delta u = 0$. For E_{H^2} , they are $\mathbf{n}^T \mathbf{H}_u \mathbf{n} = 0$.

As discussed in Section 2, many previous works in geometry processing minimize the squared Laplacian energy E_{Δ^2} subject to zero Neumann conditions ($\nabla u \cdot \mathbf{n}|_{\partial\Omega} = 0$). We obtain that the only additional natural boundary conditions are third-order: $\nabla \Delta u \cdot \mathbf{n} = 0$.

For data smoothing, zero Neumann boundary conditions encourage functions to be flat near the boundary. The null space correspondingly shrinks to contain only constant functions. If the smoothness energy dominates over the data, then a best fit constant function is found (see Fig. 3, *center*). For arbitrary domains, the zero Neumann boundary conditions create a strong boundary sensitivity (see Fig. 8).

4.5 Relationship to previous discrete energies

The null space of the squared Hessian energy E_{H^2} contains and only contains linear functions. In the context of real-time shape deformation, Wang et al. [2015] specify that this is precisely the condition needed of a smoothness energy for linear subspace design. With this property as a goal, Wang et al. and other previous works have designed discrete energies by modifying the discrete cotangent Laplace operator, $L \in \mathbb{R}^{n \times n}$, for a triangle mesh with n vertices and k edges. The unmodified discrete Laplace operator L can be derived via a finite-element discretization of the squared gradient energy in Equation (1), and the discrete energy preserves the smooth structure of having only constant functions in its null space. In contrast, the following two discrete modifications of L do not have known smooth analogs, and as such are more susceptible to errors and more challenging to analyze [Wang et al. 2017].

4.5.1 Discrete exterior calculus. Fisher et al. [2007] construct a discrete energy for tangent vector field design using discrete exterior

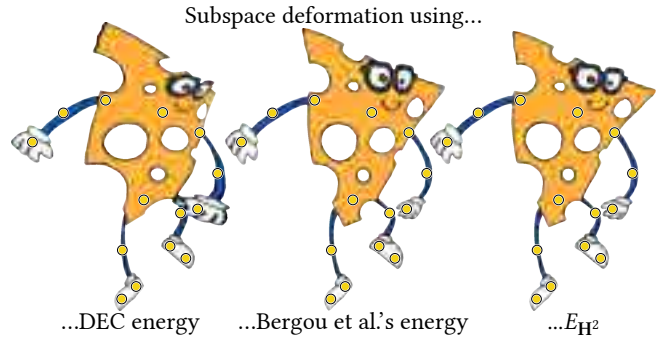


Fig. 9. The smoothness energy built using discrete exterior calculus in Equation (20) is not suitable for linear subspace design. In contrast, deformations using Equation (21) or E_{H^2} are similarly appropriate.

calculus (DEC) on triangle meshes. Their energy measures the sum of all squared discrete divergences of a given vector field \mathbf{v} via:

$$\mathbf{v}^T \star_1 d_0 \star_0^{-1} \underbrace{d_0^T \star_1 \mathbf{v}}_{\mathbf{F}}, \quad (19)$$

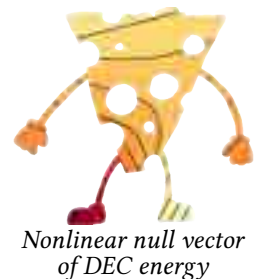
where \star_0 , \star_1 and d_0 are the discrete Hodge star operators for zero- and one-forms and the discrete differential matrix for one-forms respectively, following the adjustments when building \mathbf{F} for free boundaries as described in [Fisher et al. 2007]. If we replace the unknown vector field \mathbf{v} with the differential of an unknown scalar field $d_0 \mathbf{u}$, we can build a discrete smoothness energy of a scalar field:

$$\mathbf{u} d_0^T \underbrace{\mathbf{F}^T}_{\mathbf{M}^{-1}} \underbrace{\star_0^{-1} \mathbf{F} d_0}_{\mathbf{K}=\mathbf{L}+\mathbf{N}} \mathbf{u}, \quad (20)$$

where the matrix $\mathbf{M} \in \mathbb{R}^{n \times n}$ is the (lumped) per-vertex mass matrix, and the non-symmetric matrix $\mathbf{K} \in \mathbb{R}^{n \times n}$ matches the description of the operator in [Wang et al. 2015] constructed by adding the normal derivative operator $\mathbf{N} \in \mathbb{R}^{n \times n}$ to the standard cotangent Laplacian \mathbf{L} . This same matrix \mathbf{K} has been previously shown to contain at least linear functions in its right null space [Crane 2009].

Unfortunately, for some meshes this matrix — and by extension the energy it defines — contains other non-linear functions in its null space (see inset). This does not appear to be a simple matter of mesh resolution or numerics. From the point of view of linear subspace design (see Fig. 9), this null space is too big.

Following the DEC construction, taking the divergence of the gradient should correspond to the smooth Laplace operator ($\Delta = \nabla \cdot \nabla$). If the discrete energy in Equation (20) were structure preserving, then we would expect it to produce discrete natural boundary conditions for the squared Laplacian energy E_{Δ^2} and match its null space of all harmonic functions (see Fig. 7). However, eigen analysis of this energy for typical meshes produces a small number (≥ 3) of null modes. From a structure preservation standpoint, this null space is too small.



Convergence

L_2 difference with analytic minimizer of E_{H^2}

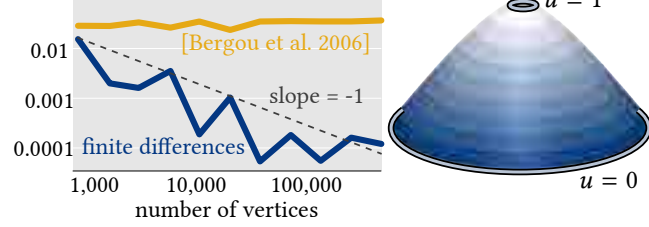


Fig. 10. On a 2D annulus, fixing inner boundary values to one and outer boundary values to zero, the analytic minimizer of E_{H^2} is a radially symmetric function with natural 3rd-order boundary conditions (see Section 4.4.1). While a finite difference discretization of E_{H^2} converges under refinement, minimizing the discrete energy in [Bergou et al. 2006] does not.

4.5.2 Non-conforming finite-elements. Wang et al. [2015] use an alternative construction to create their results [Wang et al. 2017]. Normal derivatives are added to the edge-based Laplacian $L_{cr} \in \mathbb{R}^{k \times k}$ resulting from the Crouzeix-Raviart non-conforming finite-element discretization of the squared gradient energy in Equation (1):

$$\mathbf{u}^T \mathbf{E}^T \mathbf{K}_{cr}^T \mathbf{M}_{cr}^{-1} \underbrace{\mathbf{K}_{cr}}_{L_{cr} + N_{cr}} \mathbf{E} \mathbf{u}, \quad (21)$$

where $\mathbf{E} \in \mathbb{R}^{k \times n}$ averages vertex values onto incident edges, $\mathbf{M}_{cr} \in \mathbb{R}^{k \times k}$ is the Crouzeix-Raviart mass matrix, and $N_{cr} \in \mathbb{R}^{k \times k}$ computes normal derivatives at boundary edges.

This energy corresponds to *one of two* discrete energies proposed by Bergou et al. [2006] to model plate bending under the assumption of isometric deformation as the squared Laplacian (E_{Δ^2}) of the displacement coordinate functions. Bergou et al. do not discuss boundary conditions and the discrete boundary conditions of their two proposed discrete energies *differ* (the other corresponding to E_{Δ^2} subject to zero-Neumann boundary conditions).

Perhaps surprisingly, the discrete energy in Equation (21) behaves very similarly to the squared Hessian energy E_{H^2} despite its motivation by both Bergou et al. [2006] and Wang et al. [2015] as a discrete analog to E_{Δ^2} . In contrast to the DEC energy in Equation (20), empirical tests indicate that the null space of the Bergou et al. energy contains and only contains linear functions. Resulting subspace deformations are visually indistinguishable from those using E_{H^2} (see Fig. 9).

Unfortunately, while the discrete energy in Equation (21) clearly does not behave like a discretization of E_{Δ^2} with natural boundary conditions, it also does not converge to smooth minimizers of E_{H^2} with its natural boundary conditions (see Fig. 10). It remains unclear – and outside our scope – to show whether this discrete energy converges to yet some other continuous energy.

5 DISCRETIZATION

For our experiments, we discretize the energies E_{Δ^2} and E_{H^2} using either the finite difference method on regular 2D grids restricted to a bounded subregion $\Omega \subset \mathbb{R}^2$ or the mixed finite element method on irregular triangle meshes.

5.1 Finite differences

For finite differences on a 2D regular grid, we use standard central differences to compute each entry of the Hessian at all interior nodes [Fornberg 1988] (see Appendix A).

Our discretization on grids is in agreement with those used in image processing on rectangular domains [Didas et al. 2009]. As expected, this standard finite difference discretization converges under refinement (see Fig. 10). No special treatment is required in the case of non-convex, irregular domains beyond ensuring that all interior nodes have neighboring nodes in all eight directions.

5.2 Mixed Finite Elements

The common discretization of the squared Laplacian energy with zero Neumann boundary conditions on a mesh with n vertices and m faces is constructed by *squaring* the cotangent Laplacian $L \in \mathbb{R}^{n \times n}$:

$$E_{\Delta^2}(u) \text{ subject to } \nabla u \cdot \mathbf{n} = 0|_{\partial\Omega} \approx \mathbf{u}^T \mathbf{L}^T \mathbf{M}^{-1} \mathbf{L} \mathbf{u}, \quad (22)$$

where $\mathbf{M} \in \mathbb{R}^{n \times n}$ is a (often lumped) *mass* matrix. Jacobson et al. [2010] demonstrated that this discrete energy can be derived by applying the mixed finite element method to the continuous energy minimization problem.

Following their steps, we can similarly discretize the natural boundary conditions¹ of the squared Laplacian energy:

$$E_{\Delta^2}(u) \approx \mathbf{u}^T \mathbf{L}(i, a)^T \mathbf{M}(i, i)^{-1} \mathbf{L}(i, a) \mathbf{u}, \quad (23)$$

where $\mathbf{X}(j, k)$ indicates *slicing* rows and columns corresponding to vertex lists j and k respectively, and specifically i and a are the lists of interior vertices and all vertices respectively. It is clear from this construction that *any* discrete harmonic function ($L(i, a)\mathbf{u} = 0$) will measure zero energy regardless of boundary values. This is correct from the point of view of structure preservation, and validates the arbitrary boundary behavior observed when using these natural boundary conditions for smoothing (see Fig. 3).

We can follow similar steps to use the mixed finite element method to discretize the squared Hessian energy. Beginning in the smooth setting, we introduce an auxiliary matrix-valued variable \mathbf{V} constrained to be equal to the Hessian of the unknown function on the interior of the domain via a matrix-valued Lagrange multiplier function Λ which is free on the interior and clamped to zero on the boundary:

$$\text{saddle}_{u, \mathbf{V}, \Lambda} \int_{\Omega} \left(\frac{1}{2} \mathbf{V} : \mathbf{V} + \Lambda : (\mathbf{H}_u - \mathbf{V}) \right) dA, \quad \text{s.t. } \Lambda = 0|_{\partial\Omega}.$$

Applying the Green's identity from Equation (11) to the $\Lambda : \mathbf{H}_u$ term, we *move* a derivative from u to Λ :

$$\text{saddle}_{u, \mathbf{V}, \Lambda} \int_{\Omega} \left(\frac{1}{2} \mathbf{V} : \mathbf{V} - (\nabla \cdot \Lambda) \cdot \nabla u - \Lambda : \mathbf{V} \right) dA + \quad (24)$$

$$\oint_{\partial\Omega} n^T \Lambda \nabla u \, ds. \quad \Lambda = 0|_{\partial\Omega}$$

¹Jacobson et al. [2010] mention “natural” boundary conditions for E_{Δ^2} but do not explicitly derive or discretize them. The free boundaries in their results appear to enforce zero Neumann boundary conditions.

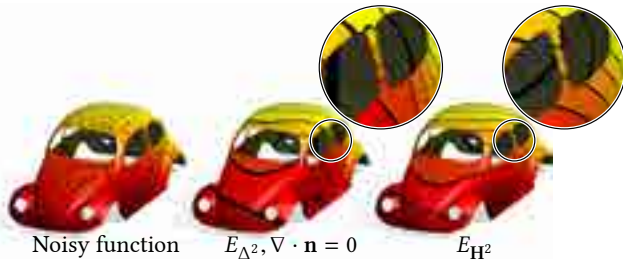


Fig. 11. Low-order boundary conditions are the *de facto* standard for smoothing in geometry processing (middle), but they introduce bias at the boundary (inset). A naïve extension of our discrete Hessian energy to 3D seems to alleviate this.

Applying a functional variation to \mathbf{V} , we immediately see that the solution must be obtained when $\mathbf{V} = \Lambda|_{\Omega}$, so we can substitute Λ out:

$$\text{saddle}_{u, \mathbf{V}} \int_{\Omega} \left(-\frac{1}{2} \mathbf{V} : \mathbf{V} - (\nabla \cdot \mathbf{V}) \cdot \nabla u \right) dA \quad (25)$$

This saddle problem involves only first derivatives. We may use standard piecewise-linear elements for u and each of the four coordinate functions of \mathbf{V} (with $\mathbf{V} = 0|_{\partial\Omega}$ because $\Lambda = 0|_{\partial\Omega}$). Factoring out the degrees of freedom corresponding to \mathbf{V} , we have a discretization of the squared Hessian energy with natural boundary conditions:

$$E_{H^2}(u) \approx \mathbf{u}^T \mathbf{G}^T \mathbf{A} \mathbf{D} \tilde{\mathbf{M}}^{-1} \mathbf{D}^T \mathbf{A} \mathbf{G} \mathbf{u}, \quad (26)$$

where $\mathbf{G}, \mathbf{A}, \mathbf{D}$, and $\tilde{\mathbf{M}}$ are the discrete gradient operator, diagonal matrix of triangle areas, discrete matrix divergence operator and discrete mass matrix (see Appendix B).

5.2.1 Alternative discretizations. Many other mixed finite-element methods have been applied to the biharmonic equation and associated energies in the past [Lamichhane 2011; Scholz 1978]. The Hellan-Herrmann-Johnson mixed formulation (see, e.g., [Comodi 1989]) has seen recent renewed interest when applied in conjunction with the discontinuous Galerkin method [Braess et al. 2017; Hoppe et al. 2016].

Conforming high-order elements are another option. For example, the Argyris element [Braess 2002] would directly allow second-order differentiation in \mathbb{R}^2 . Besides introducing a large number of degrees of freedom, it is dependent on the 2D Cartesian coordinate system, making it difficult to extend to arbitrary surfaces.

Higher-order smoothing energies have also been discretized without the finite element method. For example, in machine learning, Kim et al. [2009] discretize the Hessian energy using a nonlinear fitting approach.

5.3 Triangle meshes in 3D

So far we have only considered flat domains $\Omega \subset \mathbb{R}^2$. First experiments suggest that curved surfaces, and especially those with boundaries, might benefit analogously from our analysis of smoothness energies (see Fig. 11). We can trivially extend our mixed finite-element discretization of E_{H^2} in Section 5.2 to triangle meshes immersed in \mathbb{R}^2 by extending the gradient and matrix divergence operators in

Equation (26) to compute 3D rather than 2D vectors. This amounts to temporarily treating the Hessian as a 3×3 matrix. This extension is inspired by the construction of the discrete Laplacian for surfaces built by trivially extending the gradient operator to compute per-face 3D vectors while maintaining the property that: $\mathbf{L} = \mathbf{G}^T \mathbf{A} \mathbf{G}$. Our discretization of the Hessian and squared Hessian matrices is available open source as part of LIBGL [2017].

5.3.1 Future work: accounting for curvature.

On curved surfaces, the formulation of the Hessian energy is far more intricate. One striking difference to the planar case arises from the fact that the Hessian energy on smooth surfaces without boundary is in general no longer equal to the Laplacian energy. Indeed, integration by parts similar to Equation (10) leads to the so-called Bochner-Laplacian, which differs from the usual (Hodge-) Laplacian by a term involving Gaussian curvature. Also the naïve extension of our mixed finite element discretization of E_{H^2} to 3D leads to a different energy than the Laplacian energy (see inset). However, we do not claim this extension to be a proper discrete version of the smooth Hessian energy.



Another complication arises from the fact that even a weak formulation of the Hessian similar to Equation (24) involves covariant differentiation in the smooth setting. One future direction might be to use recent discretizations of the covariant derivative (e.g., [Azencot et al. 2015; Liu et al. 2016]) to derive a discrete Hessian energy while maintaining analogous natural boundary conditions. The original Green's identity in Equation (7) does *not* introduce curvature terms, allowing discretization of the squared Laplacian Energy using mixed FEM as used for thin shells [Bergou et al. 2006; Garg et al. 2007].

Regardless, first experiments suggest that our naïve extension to 3D might enable our main results concerning free boundaries to operate on surface meshes. We leave an analysis of our naïve extension to 3D for future work.

5.4 L_1 minimization

Instead of minimizing the L_2 -norm of the Hessian, the squared Hessian energy, we can also minimize its L_1 -norm:

$$\int_{\Omega} |\mathbf{H}_u|_F dA \quad (27)$$

This L_1 -type energy exhibits different properties from the ones of the squared Hessian energy (see Section 6).

A detailed description the discretization can be found in Appendix C. This is a simple discretization; other discretizations of the L^1 norm are discussed in works such as [Bronstein et al. 2016].

6 EXPERIMENTS & APPLICATIONS

Smoothness energies are a fundamental ingredient in geometry processing algorithms. We tour applications using the Hessian to define smoothness. We solve the resulting sparse linear systems



Fig. 12. Both weights are constrained to agree on the outer boundary. On the interior holes, E_{H^2} causes less distortion than E_{Δ^2} .

using MATLAB and quadratic programs using MOSEK [Andersen and Andersen 2000]. We verified that both our discretizations (see Section 5) converge toward the analytic minimizer of E_{H^2} on an annular domain with radially symmetric fixed value and third-order natural boundary conditions (see, e.g., Fig. 10). The computational complexity of minimizing E_{H^2} is equivalent to minimizing E_{Δ^2} : in both cases the number of non-zeros per-row is less than or equal to the size of the corresponding vertex’s two-neighborhood.

The simplest demonstration of these energies is to reconstruct a smooth function while interpolating values at specific points. For flat domains with interpolated values and normal derivatives along the boundary, the reconstructions minimizing E_{Δ^2} and E_{H^2} will agree (see Fig. 6). If the boundary is left unconstrained or partially constrained then natural boundary conditions will appear for each respective energy (see Fig. 2). In all examples, differences are best identified by observing how isolines meet with the boundary: zero Neumann boundary conditions cause the isolines to be perpendicular to the boundary, while high-order natural boundary conditions do not force this behavior (see Fig. 8).

Linear subspace design for cartoon deformation is a special case of scattered data interpolation. Instead of interpolating colors or temperatures, *displacements* are interpolated over a 2D cartoon or surface (ref., [Jacobson et al. 2014]). Minimizing E_{H^2} with its natural boundary conditions and Kronecker delta values at specified *point handles*, we can define a linear basis for smooth deformation interpolation displacements at these points. Like previous approaches [Wang et al. 2015], the minimizing functions of E_{H^2} are *linearly precise* and therefore form a form of *cage-free* generalized barycentric coordinates (“Hessian coordinates”). In contrast to previous work, like the coordinates of Wang et al., Hessian coordinates are defined by a smooth energy where it is easy to show that all and only linear functions exist in its null space (see Fig. 7). Fig. 9 shows that the deformation behavior is similar that of Wang et al. [2015]

Similarly, automatic methods for computing linear blend skinning weights have employed smoothness energies [Baran and Popović 2007]. Classic bounded biharmonic weights minimize E_{Δ^2} subject to zero Neumann boundary conditions [Jacobson et al. 2011] and bound constraints. Restricting E_{H^2} to the same value as classic bounded biharmonic weights on the outside of the shape, but not on interior holes leads to weights that preserve the characteristic shape of holes without much distortion. As can be seen in Fig. 12, bounded biharmonic weights based on E_{Δ^2} distort the interior holes

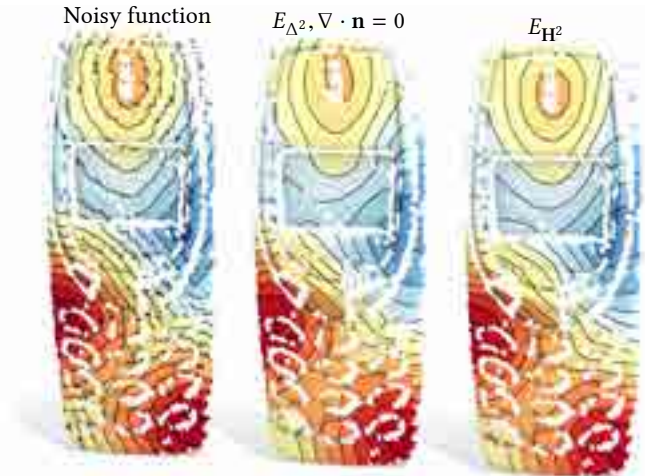


Fig. 13. Smoothing on a noisy domain. Compared to previous methods (left), the natural boundary conditions of E_{H^2} (right) better model *ignoring* the non-salient boundaries when smoothing data on a partial surface scan.

of the pants significantly under deformation, while E_{H^2} alleviates the effect.

We now turn our attention to another common use for a smoothness energy: dense data denoising or fairing. To do so, we can optimize a function u that minimizes the weighted sum of a data term (L_2 -norm of difference between u and a noisy input function f) and the L_2 smoothness energy (e.g., E_{Δ^2} or E_{H^2}). When using E_{Δ^2} for data smoothing on domains with free boundaries, previous methods (e.g., [Weinkauff et al. 2010]) enforce low-order boundary conditions to ensure that noisy boundary values are not simply interpolated (see Fig. 3).

The biasing effect of these low-order conditions is apparent: the heavier the smoothing, the more the solution becomes constant near the boundary *regardless of the data there*. In contrast, smoothing with E_{H^2} allows the solution to vary near the boundary (see Fig. 11).

In Fig. 13, we smooth a noisy simulation of heat diffusion on a range scan of a Nokia cellphone [Godil et al. 2009]. The abundance of free boundaries due to missing data highlights the effect of zero Neumann boundary conditions compared to the natural boundary conditions of E_{H^2} .

The energies E_{Δ^2} and E_{H^2} measure the L_2 -norm of the Laplacian / Hessian respectively. Minimization of such energies prefers to distribute energy smoothly over the domain. In contrast, L_1 minimization prefers to concentrate high energy at sparse locations. At




Fig. 14. Smoothing using the L_2 -norm also removes salient creases. Minimizing the L_1 -norm of the Hessian produces a piecewise-planar solution.




Fig. 15. The *Bunny* flows along L_1 -minimization of the Hessian (local energy density in pseudocolor). Energy concentrates at creases and the bunny develops smooth, nearly flat regions. This minimization achieves crease-aware smoothing of surface geometry. Minimizing the L_1 -norm of the Laplacian leads to energy concentration at *points* and a prickly appearance.

these locations we see the behavior of natural boundary conditions of the smoothness energy. Effectively, they become boundaries between low-energy regions. In Fig. 14, we smooth a toy function (a triangle wave plus noise). While L_2 smoothing E_{H^2} also rounds the peaks, L_1 smoothing E_{H^1} smooths away the noise but maintains the sharp creases.

In Fig. 16, we smooth the noisy height data of a cathedral rooftop. Minimizing the L_1 -norm of the Laplacian (E_{Δ^1}) concentrates energy at isolated points, producing a circus tent appearance. In contrast, minimizing the L_1 -norm of the Hessian (E_{H^1}) concentrates energy along creases, producing a piecewise planar rooftop.

Under the L_1 -norm, the difference between minimizers does not rely on the presence of a boundary. Indeed, even for closed surfaces with *no* boundary, we see very different behavior. In Fig. 15, we treat the surface's geometry as the input data f and take smoothing steps where the data-versus-smoothness weight acts as an implicit time-step parameter controlling a geometric flow. The E_{H^1} flow of the *Bunny* forms 1D creases bounding smooth, near-developable (low Gaussian curvature) patches. This application is inspired by image smoothing with sparse norms [Xu et al. 2011] and as such the results are reminiscent of surface contrast enhancing methods (e.g., [He and Schaefer 2013]). In contrast, the E_{Δ^1} flow quickly concentrates energy at isolated points suspended in a near-minimal (low mean curvature) surface.

7 CONCLUSION & FUTURE WORK

Energies built with the Hessian rather than the Laplacian unlock high-order boundary conditions that are especially useful when boundaries are to be conceptually ignored during problem modeling.

In future work, we would like to investigate formulations of E_{H^2} for curved surfaces that account for the covariant derivative, as well as corresponding discretizations (see Section 5.3.1). For flat domains, it would be interesting to explore boundary-only discretizations (cf. [Chen and Weber 2015]) or subspace deformation for solid objects. Initial derivations suggest that point constraints may not lead to smooth minimizers of second-order smoothness energies for solids. One possibility, however daunting, may be to consider the squared Frobenius norm of the three-tensor of third derivatives, an energy with only quadratic functions in its null space.

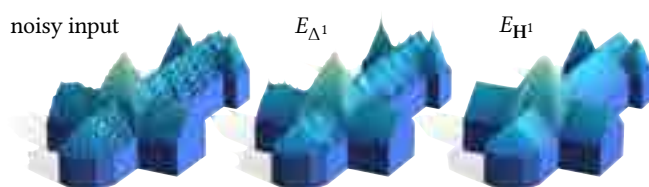


Fig. 16. Smoothing a cathedral roof. Unlike L_2 minimization, the behavior of $E_{\Delta^1} = \int |\Delta u| dA$ and $E_{H^1} = \int |\mathbf{H}_u|_F dA$ are dramatically different in the interior.

We hope that our work sheds new light on familiar problems and provides insights into the power of the natural boundary conditions of the squared Hessian energy for geometry processing problems. Many applications that are currently using the squared Laplacian energy with zero Neumann boundary conditions can potentially profit by trying the squared Hessian energy with natural boundary conditions alongside it.

ACKNOWLEDGEMENTS

We would like to thank Ladislav Kavan, Denis Zorin, and Jinshuo Dong for early discussions; Yu Wang for sharing source code; Paul Kry for hosting the 2017 Bellairs workshop; and organizers of the Dagstuhl Seminar #17232.

This work is funded in part by the Binational Science Foundation (US-Israel) Award 2012376, the National Science Foundation Award IIS-14-09286, the NSERC Discovery Grants (RGPIN-2017-05235 & RGPAS-2017-507938), a Canada Research Chair award, the Connaught Fund, and a gift from Adobe Systems Inc.

REFERENCES

- E. D. Andersen and K. D. Andersen. 2000. The MOSEK interior point optimizer for linear programming: an implementation of the homogeneous algorithm. In *High Performance Optimization*. Kluwer Academic Publishers, 197–232.
- James Andrews, Pushkar Joshi, and Nathan Carr. 2011. A Linear Variational System for Modeling From Curves. *j-CGF* (2011).
- Omri Azencot, Maks Ovsjanikov, Frédéric Chazal, and Mirela Ben-Chen. 2015. Discrete Derivatives of Vector Fields on Surfaces – An Operator Approach. *ACM Trans. Graph.* 34, 3, Article 29 (May 2015), 13 pages. <https://doi.org/10.1145/2723158>
- Ilya Baran and Jovan Popović. 2007. Automatic Rigging and Animation of 3D Characters. *ACM Trans. Graph.* 26, 3 (2007), 72:1–72:8.
- Christopher Batty. 2010. *Simulating viscous incompressible fluids with embedded boundary finite difference methods*. Ph.D. Dissertation. University of British Columbia.

- M. Belkin and P. Niyogi. 2004. Semi-supervised Learning on Riemannian Manifolds. *Machine Learning* 56, 1-3 (2004), 209–239.
- Miklos Bergou, Max Wardetzky, David Harmon, Denis Zorin, and Eitan Grinspun. 2006. A Quadratic Bending Model for Inextensible Surfaces. In *Proc. SGP*. 227–230.
- Morten Bojsen-Hansen and Chris Wojtan. 2016. Generalized Non-Reflecting Boundaries for Fluid Re-Simulation. *ACM Trans. Graph.* 35, 4 (2016).
- F. L. Bookstein. 1989. Principal Warps: Thin-Plate Splines and the Decomposition of Deformations. *IEEE Trans. Pattern Anal. Mach. Intell.* (1989).
- Mario Botsch and Leif Kobbelt. 2004. A Remeshing Approach to Multiresolution Modeling. In *Proc. SGP*. 189–196.
- M. Botsch and L. Kobbelt. 2005. Real-time shape editing using radial basis functions. *Comp. Graph. Forum* 24, 3 (2005), 611–621.
- D. Braess. 2002. *Finite Elements: Theory, fast solvers and applications in solid mechanics, 2nd edn.* Vol. 13, 1500 pages.
- D. Braess, A. S. Pechstein, and J. Schöberl. 2017. An Equilibration Based A Posteriori Error Estimate for the Biharmonic Equation and Two Finite Element Methods. *ArXiv* (2017). arXiv:1705.07607
- A. Bronstein, Y. Choukroun, R. Kimmel, and M. Sela. 2016. Consistent Discretization and Minimization of the L1 Norm on Manifolds. In *Fourth International Conference on 3D Vision (3DV)*.
- Chen Cao, Derek Bradley, Kun Zhou, and Thabo Beeler. 2015. Real-time High-fidelity Facial Performance Capture. *ACM Trans. Graph.* (2015).
- Chen Cao, Yanlin Weng, Shun Zhou, Yiyong Tong, and Kun Zhou. 2014. Facewarehouse: A 3d facial expression database for visual computing. *IEEE TVCG* (2014).
- Renjie Chen and Ofir Weber. 2015. Bounded Distortion Harmonic Mappings in the Plane. *ACM Trans. Graph.* (2015).
- Daniel Cohen-Steiner and Mathieu Desbrun. 2002. *Hindsight: LSCM and DNCP are one and the same.* Technical Report, INRIA-USC.
- MI Comodi. 1989. The Hellan-Herrmann-Johnson method: some new error estimates and postprocessing. *mathematics of computation* 52, 185 (1989), 17–29.
- Richard Courant and David Hilbert. 1924. *Methoden der Mathematischen Physik I.* Keenan Crane. 2009. *A Cotangent Laplacian for Surfaces With Boundary.* Tech. Report.
- Fernando de Goes, Beibei Liu, Max Budninskiy, Yiyong Tong, and Mathieu Desbrun. 2014. Discrete 2-Tensor Fields on Triangulations. *Comp. Graph. Forum* (2014).
- Mathieu Desbrun, Mark Meyer, and Pierre Alliez. 2002. Intrinsic Parameterizations of Surface Meshes. *Comp. Graph. Forum* 21, 3 (2002), 209–218.
- Mathieu Desbrun, Mark Meyer, Peter Schröder, and Alan H. Barr. 1999. Implicit fairing of irregular meshes using diffusion and curvature flow. In *Proc. of ACM SIGGRAPH*.
- Stephan Didas and Joachim Weickert. 2004. Higher order variational methods for noise removal in signals and images. *Saarbrücken, Saarland Universitesi* (2004).
- Stephan Didas, Joachim Weickert, and Bernhard Burgeth. 2009. Properties of higher order nonlinear diffusion filtering. *Journal of mathematical imaging and vision* 35, 3 (2009), 208–226.
- David L. Donoho and Carrie Grimes. 2003. Hessian eigenmaps: Locally linear embedding techniques for high-dimensional data. *Proc. of the National Academy of Sciences* (2003).
- Lawrence C. Evans. 1998. *Partial differential equations.*
- Carlos A. Felippa. 2017. Advanced Finite Element Methods - Course Notes - <https://www.colorado.edu/engineering/CAS/courses/d/AFEM.d/>. (April 2017).
- Mark Finch, John Snyder, and Hugues Hoppe. 2011. Freeform vector graphics with controlled thin-plate splines. *ACM Trans. Graph.* (2011).
- M. Fisher, P. Schröder, M. Desbrun, and H. Hoppe. 2007. Design of tangent vector fields. *ACM Trans. Graph.* 26, 3 (2007).
- Bengt Fornberg. 1988. Generation of finite difference formulas on arbitrarily spaced grids. *Mathematics of computation* (1988).
- A. Garg, E. Grinspun, M. Wardetzky, and D. Zorin. 2007. Cubic shells. *Proceedings of the 2007 ACM SIGGRAPH/Eurographics symposium on Computer animation* (2007), 91–98.
- I.M. Gelfand and S.V. Fomin. 1963. *Calculus of Variations.*
- T.G. Georgiev. 2004. Photoshop Healing Brush: a Tool for Seamless Cloning. In *Proc. ECCV ACV Workshop*. 1–8.
- Mariano Giaquinta and Stefan Hildebrandt. 1996. *Calculus of Variations I.*
- Y.I. Gingold, P. Davidson, J. Han, and D. Zorin. 2006. A Direct Texture Placement and Editing Interface. In *Proc. UIST*.
- A Godil, A Axenopoulos, P Daras, T Furuya, and R Ohbuchi. 2009. SHREC 2009: shape retrieval contest of partial 3D models. In *Proc. Eurographics Workshop 3D Object Retrieval*.
- Lei He and Scott Schaefer. 2013. Mesh Denoising via L0 Minimization. *ACM Trans. Graph.* (2013).
- Klaus Hildebrandt, Christian Schulz, Christoph Von Tycowicz, and Konrad Polthier. 2011. Interactive surface modeling using modal analysis. *ACM Trans. Graph.* (2011).
- Ronald Hoppe, Dietrich Braess, and Christopher Linsenmann. 2016. A two-energies principle for the biharmonic equation and an a posteriori error estimator for an interior penalty discontinuous Galerkin approximation. (11 2016).
- Jin Huang, Xiaohan Shi, Xinguo Liu, Kun Zhou, Li-Yi Wei, Shang-Hua Teng, Hujun Bao, Baining Guo, and Heung-Yeung Shum. 2006. Subspace gradient domain mesh deformation. *ACM Trans. Graph.* 25, 3 (2006), 1126–1134.
- Alec Jacobson, Ilya Baran, Jovan Popović, and Olga Sorkine. 2011. Bounded Biharmonic Weights for Real-Time Deformation. *ACM Trans. Graph.* 30, 4 (2011).
- Alec Jacobson, Zhigang Deng, Ladislav Kavan, and JP Lewis. 2014. Skinning: Real-time Shape Deformation. In *ACM SIGGRAPH 2014 Courses*.
- Alec Jacobson, Daniele Panozzo, et al. 2017. libigl: A simple C++ geometry processing library. (2017). <http://libigl.github.io/libigl/>.
- Alec Jacobson, Elif Tosun, Olga Sorkine, and Denis Zorin. 2010. Mixed Finite Elements for Variational Surface Modeling. In *Proc. SGP*.
- Alec Jacobson, Tino Weinkauff, and Olga Sorkine. 2012. Smooth Shape-Aware Functions with Controlled Extrema. In *Proc. SGP*.
- Ben Jones, Nils Thuerey, Tamar Shinar, and Adam W. Bargteil. 2016. Example-based Plastic Deformation of Rigid Bodies. *ACM Trans. Graph.* (2016).
- Ladislav Kavan, Dan Gerszewski, Adam Bargteil, and Peter-Pike Sloan. 2011. Physics-inspired Upsampling for Cloth Simulation in Games. *ACM Trans. Graph.* (2011).
- Kwang I. Kim, Florian Steinke, and Matthias Hein. 2009. Semi-supervised Regression using Hessian energy with an application to semi-supervised dimensionality reduction. In *Advances in Neural Information Processing Systems*.
- Bishnu P. Lamichhane. 2011. A stabilized mixed finite element method for the biharmonic equation based on biorthogonal systems. *Journal of Comp. and Appl. Math.* (2011).
- E. Landreneau and S. Schaefer. 2010. Poisson-Based Weight Reduction of Animated Meshes. *Comp. Graph. Forum* 29, 6 (2010), 1945–1954.
- Stamatis Lefkimmiatis, Aurélien Bourquard, and Michael Unser. 2012. Hessian-based norm regularization for image restoration with biomedical applications. *IEEE Transactions on Image Processing* 21, 3 (2012), 983–995.
- B. Lévy, S. Petitjean, N. Ray, and J. Maillot. 2002. Least squares conformal maps for automatic texture atlas generation. *ACM Trans. Graph.* 21, 3 (2002), 362–371.
- W.L. Li. 2000. Free vibrations of beams with general boundary conditions. *Journal of Sound and Vibration* 237 (2000).
- Yaron Lipman, Raif Rustamov, and Thomas Funkhouser. 2010. Biharmonic Distance. *ACM Trans. Graph.* 29, 3 (2010).
- Beibei Liu, Yiyong Tong, Fernando De Goes, and Mathieu Desbrun. 2016. Discrete Connection and Covariant Derivative for Vector Field Analysis and Design. *ACM Trans. Graph.* (2016).
- XY Liu, C-H Lai, and Kyriacos A Pericleous. 2015. A fourth-order partial differential equation denoising model with an adaptive relaxation method. *International Journal of Computer Mathematics* 92, 3 (2015), 608–622.
- Marius Lysaker, Arvid Lundervold, and Xue-Cheng Tai. 2003. Noise removal using fourth-order partial differential equation with applications to medical magnetic resonance images in space and time. *IEEE Trans. on image processing* (2003).
- Marius Lysaker and Xue-Cheng Tai. 2006. Iterative image restoration combining total variation minimization and a second-order functional. *IJCV* (2006).
- Patrick Mullen, Yiyong Tong, Pierre Alliez, and Mathieu Desbrun. 2008. Spectral Conformal Parameterization. In *Proc. SGP*.
- Ulrich Pinkall and Konrad Polthier. 1993. Computing discrete minimal surfaces and their conjugates. *Experiment. Math.* 2, 1 (1993), 15–36.
- Jean Marc Roth. 2009. *Higher Order Anisotropic Smoothing of Images.* Ph.D. Dissertation. Saarland University.
- Leonid I Rudin, Stanley Osher, and Emad Fatemi. 1992. Nonlinear total variation based noise removal algorithms. *Physica D: Nonlinear Phenomena* 60, 1 (1992), 259–268.
- Raif M Rustamov. 2011. Multiscale biharmonic kernels. (2011).
- R. Scholz. 1978. A mixed method for 4th order problems using linear finite elements. *RAIRO Anal. Numér.* (1978).
- Olga Sorkine, Yaron Lipman, Daniel Cohen-Or, Marc Alexa, Christian Rössl, and Hans-Peter Seidel. 2004. Laplacian Surface Editing. In *Proc. SGP*. 179–188.
- Boris Springborn, Peter Schröder, and Ulrich Pinkall. 2008. Conformal Equivalence of Triangle Meshes. *ACM Trans. Graph.* (2008).
- Gabriele Steidl. 2006. A note on the dual treatment of higher-order regularization functionals. *Computing* 76, 1-2 (2006), 135–148.
- Gabriele Steidl, Stephan Didas, and Julia Neumann. 2005. Relations between higher order TV regularization and support vector regression. In *International Conference on Scale-Space Theories in Computer Vision*. Springer, 515–527.
- Florian Steinke and Matthias Hein. 2009. Non-parametric regression between manifolds. In *Advances in Neural Information Processing Systems*. 1561–1568.
- Daniel Šykora, Ladislav Kavan, Martin Čadík, Ondřej Jamriska, Alec Jacobson, Brian Whited, Maryann Simmons, and Olga Sorkine-Hornung. 2014. Ink-and-Ray: Bas-Relief Meshes for Adding Global Illumination Effects to Hand-Drawn Characters. *ACM Trans. Graph.* (2014).
- Demetri Terzopoulos. 1984. Multilevel reconstruction of visual surfaces: variational principles and finite-element representations. In *Multiresolution image processing and analysis*. Springer, 237–310.
- Demetri Terzopoulos. 1988. The computation of visible-surface representations. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 10, 4 (1988), 417–438.
- E. Tosun, Y. I. Gingold, J. Reisman, and D. Zorin. 2007. Shape optimization using reflection lines. In *Proc. SGP*.

- Yu Wang, Alec Jacobson, Jernej Barbič, and Ladislav Kavan. 2015. Linear Subspace Design for Real-time Shape Deformation. *ACM Trans. Graph.* (2015).
- Yu Wang, Alec Jacobson, Jernej Barbič, and Ladislav Kavan. 2017. *Error in "Linear Subspace Design for Real-Time Shape Deformation"*. Tech. Report.
- Ofir Weber, Roi Poranne, and Craig Gotsman. 2012. Biharmonic coordinates. In *Computer Graphics Forum*, Vol. 31. Wiley Online Library, 2409–2422.
- Tino Weinkauff, Yotam Gingold, and Olga Sorkine. 2010. Topology-based Smoothing of 2D Scalar Fields with C^1 -Continuity. *Comput. Graph. Forum* 29, 3 (June 2010), 1221–1230.
- Li Xu, Cewu Lu, Yi Xu, and Jiaya Jia. 2011. Image Smoothing via L0 Gradient Minimization. *ACM Trans. Graph.* (2011).
- Y-L You and Mostafa Kaveh. 2000. Fourth-order partial differential equations for noise removal. *IEEE Transactions on Image Processing* 9, 10 (2000), 1723–1730.
- Jing Yuan, Christoph Schnörr, and Gabriele Steidl. 2009. Total-variation based piecewise affine regularization. In *International Conference on Scale Space and Variational Methods in Computer Vision*. Springer, 552–564.
- Hao Zhang, Oliver van Kaick, and Ramsay Dyer. 2007. Spectral Methods for Mesh Processing and Analysis. In *Proc. of Eurographics State-of-the-art Report*. 1–22.
- Kun Zhou, Jin Huang, John Snyder, Xinguo Liu, Hujun Bao, Baining Guo, and Heung-Yeung Shum. 2005. Large mesh deformation using the volumetric graph Laplacian. *ACM Trans. Graph.* 24, 3 (2005), 496–503.
- Kun Zhou, Weiwei Xu, Yiyang Tong, and Mathieu Desbrun. 2010. Deformation Transfer to Multi-Component Objects. *Comp. Graph. Forum* 29, 2 (2010).

A FINITE DIFFERENCES

In order to discretize E_{H^2} over a domain Ω embedded in a grid with n total nodes and \hat{n} interior nodes, we define a sparse matrix $\mathbf{H} = \left[\mathbf{H}_{xx}^T \ \mathbf{H}_{yy}^T \ \sqrt{2}\mathbf{H}_{xy}^T \right]^T \in \mathbb{R}^{3\hat{n} \times n}$ so that each row of \mathbf{H}_{xx} approximates the second derivative of the unknown function at the corresponding interior grid node: $\mathbf{H}_{xx}^i \mathbf{u} \approx \partial^2 u(\mathbf{x}_i) / \partial x^2$ and analogously for rows in \mathbf{H}_{yy} and \mathbf{H}_{xy} .

We use standard, second-order stencils for each term, i.e.,:

$$\begin{aligned} \mathbf{H}_{xx}^i \mathbf{u} &= \frac{u_{i,j-1} - 2u_{i,j} + u_{i,j+1}}{h^2} & (28) \\ \mathbf{H}_{yy}^i \mathbf{u} &= \frac{u_{i-1,j} - 2u_{i,j} + u_{i+1,j}}{h^2} \\ \mathbf{H}_{xy}^i \mathbf{u} &= \mathbf{H}_{yx}^i \mathbf{u} = \frac{u_{i-1,j-1} + u_{i-1,j+1} - u_{i+1,j-1} - u_{i+1,j+1}}{4h^2}. \end{aligned}$$

B MIXED FINITE ELEMENTS

Solving the piecewise-linear discretization of Equation (25) by differentiating with respect to all degrees of freedom, we have a system of linear equations in matrix form:

$$\begin{pmatrix} \tilde{\mathbf{M}} & \mathbf{D}^T \mathbf{A} \mathbf{G} \\ \mathbf{G}^T \mathbf{A} \mathbf{D} & 0 \end{pmatrix} \begin{pmatrix} \mathbf{V}_{xx} \\ \mathbf{V}_{xy} \\ \mathbf{V}_{yx} \\ \mathbf{V}_{yy} \\ \mathbf{u} \end{pmatrix} = 0 \quad (29)$$

where $\tilde{\mathbf{M}} \in \mathbb{R}^{4|i| \times 4|i|}$ repeats the mass matrix $\mathbf{M}(i, i)$ along the diagonal, $\mathbf{A} \in \mathbb{R}^{2m \times 2m}$ is a diagonal matrix containing triangle areas, and $\mathbf{D} \in \mathbb{R}^{2m \times 4|i|}$ computes the matrix divergence defined by:

$$\mathbf{D} = \begin{pmatrix} \mathbf{G}(x, i) & \mathbf{G}(y, i) & 0 & 0 \\ 0 & 0 & \mathbf{G}(x, i) & \mathbf{G}(y, i) \end{pmatrix}, \quad (30)$$

where $\mathbf{G} \in \mathbb{R}^{2m \times n}$ is the usual gradient operator and $\mathbf{G}(x, i)$ selects the rows and columns corresponding to the x -components and interior vertices respectively. Finally, if we use lumped mass matrices we can efficiently condense this system and use it to define a

discretization of our original energy:

$$E_{H^2}(u) \approx \mathbf{u}^T \mathbf{G}^T \mathbf{A} \tilde{\mathbf{D}} \mathbf{M}^{-1} \mathbf{D}^T \mathbf{A} \mathbf{G} \mathbf{u}. \quad (31)$$

C L_1 MINIMIZATION

In the smooth setting, the L^1 Frobenius norm of the Hessian is:

$$\int_{\Omega} |\mathbf{H}_u|_F \, dA \quad (32)$$

This can be minimized by introducing an auxiliary matrix-valued variable equal to the element-wise absolute value of the Hessian $\mathbf{Y} = |\mathbf{H}_u|$ and solving the constrained optimization problem:

$$\min_{\mathbf{u}, \mathbf{Y}} \int_{\Omega} \mathbf{1}^T \mathbf{Y} \mathbf{1} \, dA \quad (33)$$

$$\text{subject to } \mathbf{H}_u \leq \mathbf{Y}, \mathbf{H}_u \geq -\mathbf{Y}, \text{ and } \mathbf{Y} \geq 0, \quad (34)$$

where $\mathbf{1}$ is a vector of ones (with appropriate length).

Using our discrete matrices, this becomes a linear program

$$\min_{\mathbf{u}, \mathbf{y}} \mathbf{1}^T \tilde{\mathbf{M}} \mathbf{y} \quad (35)$$

$$\text{subject to } \mathbf{H} \mathbf{u} \geq -\mathbf{y} \quad (36)$$

$$\mathbf{H} \mathbf{u} \leq \mathbf{y} \quad (37)$$

$$\mathbf{y} \geq 0, \quad (38)$$

where $\mathbf{y} \in \mathbb{R}^{4\hat{n}}$ is a vectorization of per-vertex 2×2 matrices on a mesh with \hat{n} interior vertices and $\mathbf{H} = \mathbf{D}^T \mathbf{A} \mathbf{G}$.

When combined with other quadratic energies such as a data-term, this transforms into a quadratic program. We solve such problems with MOSEK [Andersen and Andersen 2000].

Received October 2017; final version December 2018; accepted February 2018