Painless, highly accurate discretizations of the Laplacian on a smooth manifold

Laurent Demanet
Department of Mathematics, Stanford University

December 2006

Abstract

This note presents an elementary numerical scheme for discretizing the Laplace-Beltrami operator on smooth manifolds, as a difference operator of arbitrary high order. The method is coordinate-free, and involves very little geometry.

1 Introduction

This note is a proof of concept that the Laplace-Beltrami operator—or Laplacian for short—on a smooth manifold can be discretized in a truly accurate manner, without the trouble of going through curvilinear coordinates, patches, and other niceties of Riemannian geometry. Good schemes for the Laplacian are in turn useful for accurate simulation of evolution (partial differential) equations *on* the manifold; or evolution equations *for* the manifold; or computation of Laplacian eigenfunctions; or definition of diffusion maps and diffusion wavelets [7, 8, 9]. Notable applications include mesh enhancement in computer graphics [12, 15], and high-dimensional classification, clustering, learning, and dimensionality reduction [1, 8, 10, 14].

Content. Let us assume that our manifold $\mathcal{M}$ is C^∞ and embedded in $\mathbb{R}^3$ (although embedding in $\mathbb{R}^n$, $n > 3$, would pose little difficulty.) The main idea of this paper is that the Laplace-Beltrami operator $\Delta_{\mathcal{M}}$ acting on a function defined on the manifold $\mathcal{M}$ is equal to the standard Laplacian acting on the *normal extension* of the function in $\mathbb{R}^3$, in a strip around $\mathcal{M}$. The normal extension of a function outside $\mathcal{M}$ is defined as constant along normals to $\mathcal{M}$. This property of the Laplace-Beltrami operator can then be leveraged to define an accurate difference operator in the tangent plane to each vertex of a triangulation, or point cloud. The main result of this note is a theorem that guarantees arbitrary high order of convergence for the discrete Laplace-Beltrami, under a mild nondegeneracy assumption of the mesh, related to solvability of the interpolation problem.

Related work. Low-order difference formulas for the Laplace-Beltrami operator are well-known; see for instance [11, 12, 15] and references therein. For convergence for difference formulas on random point clouds, see [2, 5]. High-order formulas for the Euclidean Laplacian are also well-known, see [6]. The finite-element method can also produce high order Euclidean discretizations, of course, although the construction of smooth basis functions may be tedious. As for computation of normals to manifolds, and other differential quantities, see for instance [4, 3].

Acknowledgments. I would like to thank Peter Schröder, Mathieu Desbrun, Mikhail Belkin and Facundo Memoli for interesting discussions and encouragements.

2 Preliminaries

We focus our discussion to 2-dimensional manifolds for simplicity but the situation is algorithmically very similar in higher dimension. It does not matter whether the manifold is compact or not, since the Laplacian is a local operator. Functions on the manifold are sampled at the vertices x_i of a point cloud, made into a mesh: each point x_i has immediate neighbors x_j to which it is connected. The mesh diameter h is the longest edge linking two immediate neighbors. In what follows we will consider neighborhoods which may be larger than the immediate neighborhood, but still of diameter $O(h)$.

The material in this section is presented for the convenience of the reader, and is not new.

2.1 Discrete Euclidean Laplacians

As a background, let us first address the problem of discretizing the Laplacian in the Euclidean case. Take a smooth function f for which we know the samples $f(x_i)$ at points $\{x_i\}$ in $\mathbb{R}^2$, with the mesh structure explained above. For every point x_i we wish to find a set of scalar coefficients w_{ij}^{Δ}, nonzero in a reasonably small neighborhood of x_i, so that

$$\Delta_i f = \sum_{x_j \text{ close to } x_i} w_{ij}^{\Delta} f(x_j)$$

would be a good approximation to the true Laplacian $\Delta f(x_i)$.

A straightforward way of obtaining the weights w_{ij}^{Δ} is to match the Taylor expansions of $\Delta_i f$ to $\Delta f(x_i)$ at every point x_i. We summarize the result:

Lemma 1. *Let p be a positive integer and $f \in C^{p+2}(\mathcal{M})$. For each x_i, choose a neighborhood $\{x_j : j \sim i\}$ (including x_i) and denote $j \sim i$ when x_j is a neighbor of x_i. Let h be the diameter of the mesh $\{x_i\}$. Then the following two statements are equivalent.*

1. *$\Delta_i f = \sum_{j\sim i} w_{ij}^{\Delta} f(x_j)$ is a pointwise convergent discretization of order $p \geq 1$ to $\Delta f(x_i)$. This means that for some constant $C_p > 0$,*

$$\sup_i |\Delta_i f - \Delta f(x_i)| \leq C_p \cdot \max_{j\sim i}(h^2 |w_{ij}|) \cdot h^p \sup_x \|\nabla^{p+2} f(x)\|. \tag{1}$$

   The notation $\nabla^{p+2} f$ refers to all the partial derivatives of f of order $p+2$.

2. *The weights w_{ij}^{Δ} obey the following relations, valid for every i. We have put $e_{ij} = x_j - x_i$.*

$$\begin{aligned}
\sum_{j\sim i} w_{ij}^{\Delta} &= 0, \\
\sum_{j\sim i} w_{ij}^{\Delta} e_{ij} &= 0, \\
\sum_{j\sim i} w_{ij}^{\Delta} e_{ij} \otimes e_{ij} &= \textit{Id. + Antisymm.} \\
\sum_{j\sim i} w_{ij}^{\Delta} e_{ij} \otimes e_{ij} \otimes e_{ij} &= 0 \\
&\dots \\
\sum_{j\sim i} w_{ij}^{\Delta} e_{ij} \otimes \dots \otimes e_{ij} &= 0.
\end{aligned}$$

   The symbol $\otimes$ denotes the tensor product, in coordinates $(u \otimes v)_{kl} = u_k v_l$. In the last equality there are $p+1$ factors e_{ij}.

The proof of this lemma is based on elementary Taylor expansions and left to the reader.

If the linear system consisting of the $p+2$ (tensor) equations written above can be solved for, then any solution gives weights w_{ij}^{Δ} that correspond to a discretization of the Laplacian of order p. Formulating practical, sufficient conditions for solvability is a difficult problem that this note has no ambition of solving; let us only remark that $\Delta_i f$ can be interpreted as the Laplacian of an interpolant, therefore a sufficient condition for solvability of the system for w_{ij}^{Δ} is uniqueness of the multivariate interpolation problem of order $p+1$. In a nutshell, two conditions need to be met for interpolation: 1) each neighborhood $\{x_j : j \sim i\}$ needs to contain enough points, and 2) each neighborhood $\{x_j : j \sim i\}$ needs to be *poised*, in the sense that no algebraic curve of degree $p+1$ passes through the points[1]. An algebraic curve of order p is the zero level set of a (nonzero) bivariate polynomial of order p. See [13] for details on the interpolation problem.[2]

For practical computations, the third equation for w_{ij}^{Δ} in Lemma 1 should be rewritten as $\sum_j w_{ij}^{\Delta} e_{ij} \otimes_S e_{ij} = \mathrm{Id.}$, where $\otimes_S$ is the symmetric tensor product, $(u \otimes_S v)_{kl} = u_k v_l + u_l v_k$. The system obtained is roughly Vandermonde, and in case the reader would find black box solvers inadequate, there is an elegant inversion method based on a block-LU factorization, as described in [13].

In what follows, we also require that the conditioning of the weights' computation does not deteriorate as the mesh diameter decreases. We can gather the requirements in the following definition.

[1] Note that sets of points chosen at random from any non-degenerate probability distribution are poised with probability 1.

[2] Note also that uniqueness of interpolation of order $p+1$ is not necessary: fewer points may be needed for computation of the Laplacian than for interpolation. The criterion becomes: $\{x_j : j \sim i\}$ is poised if the points do not belong to the zero level set of a nonzero bivariate polynomial of order $p+1$ *whose harmonic part has been removed.*

Definition 1. *Fix $p > 0$. Consider a sequence of meshes $\{x_i\}$, with diameter $h \to 0$. Let M_h be the matrix for the linear system in Lemma 1. Then the sequence of meshes is* poised for the Laplacian *at order p if there exists a left inverse M_h^{-1} for M_h such that $\|M_h^{-1}\| \leq C \cdot h^{-2}$, where C may depend on p (and on the choice of norm), but not on h.*

If a sequence of meshes is poised, then the factor $\max_{j\sim i}(h^2|w_{ij}|)$ can be removed in equation (1). Definition 1 generalizes in a straightforward manner to sequences of collections of neighborhoods, instead of sequence of meshes.

2.2 Computation of the normals

A second preliminary concern is computation of the *normal* to $\mathcal{M}$, at each point x_i, to arbitrary high order of accuracy.

Let us start by observing that, near every point x_i of the smooth manifold $\mathcal{M}$, assumed of class C^∞, there always exist coordinates (y_1, y_2, z) of $\mathbb{R}^3$ in which $x_i = (0, 0, 0)$, and $\mathcal{M}$ is locally the graph of a C^∞ function $z = f(y_1, y_2)$. If such a representation is available, then the normal n_i at x_i is given by

$$n_i = \frac{(\nabla f(0,0), -1)^t}{\sqrt{|\nabla f(0,0)|^2 + 1}}.$$

In that case, we are led to the problem of computing the gradient of a function from its samples $z_j = f(y_{1j}, y_{2j})$, where $x_j = (y_{1j}, y_{2j}, z_j)$.

Let us first explain why we claim that there exist coordinates for which $\mathcal{M}$ is the graph of a function. Of course taking the two independent variables y_1, y_2 in the tangent plane at x, and the remaining coordinate z along the normal at x would be a preferred orthobasis for this purpose, but small perturbations about this canonical choice will work just as well. In fact, an application of the implicit function theorem would show that $O(1)$ perturbations of the angles about tangency and normalcy are allowed, and still yield a valid frame in which $\mathcal{M}$ is locally the graph of a function.

Therefore an approximate, possibly not very accurate choice of normal n_i^{poor} at a point x_i can be used to define y_1, y_2 in the "poor tangent" plane, and z along the "poor normal", such that $\mathcal{M}$ is locally defined by some function $z = f(y_1, y_2)$. In practice, the poor normal can be chosen as the normal vector to any triangle formed by x_i and two of its immediate neighbors x_j. The accuracy of this choice of normal is $O(h)$ where h is the mesh diameter: the validity of the graph representation is therefore clear for h small enough. (If h is large we make no claim of accuracy anyway.) An obvious scaling argument shows that there will be $O(h^{-2})$ neighbors x_j of x_i for which the graph representation is valid.

When coordinates have been chosen, we can repeat the reasoning in the previous section to approximate the gradient of a function as

$$\nabla_i f = \sum_{x_j \text{ close to } x_i} w_{ij}^\nabla f(x_j).$$

Now, each w_{ij}^∇ is a vector with 2 components. The result is as follows.

Lemma 2. *Let q be a positive integer and $f \in C^{q+1}(\mathcal{M})$. For each x_i, choose a neighborhood $\{x_j : j \sim i\}$(including x_i), where $j \sim i$ when x_j is a neighbor of x_i. Then the following two statements are equivalent.*

1. *$\nabla_i f$ is a pointwise convergent discretization of order $q \geq 1$ to $\nabla f(x_i)$. This means that for some constant $C > 0$,*

$$\sup_{x_i} \|\nabla_i f - \nabla f(x_i)\| \leq C \cdot \max(h^{-1}|w_{ij}^\nabla|)| \cdot h^q \sup_x \|\nabla^{q+1} f(x)\|. \tag{2}$$

2. *The weights w_{ij}^∇ obey the following relations, valid for every i. We have put $e_{ij} = x_j - x_i$.*

$$\sum_j w_{ij}^\nabla = 0,$$

$$\sum_j w_{ij}^\nabla \otimes e_{ij} = Id,$$

$$\sum_j w_{ij}^\nabla \otimes e_{ij} \otimes e_{ij} = 0,$$

$$\dots$$

$$\sum_j w_{ij}^\nabla \otimes e_{ij} \otimes \dots \otimes e_{ij} = 0.$$

In the last equality there are q factors e_{ij}.

Note that the function f of which the gradient needs to be computed is in this section an explicit, local representation of the manifold as a height field, and not the function interpolating the data as in Section 2.1.

The solvability discussion is the similar to that in the previous section. In an obvious generalization, we define a sequence of meshes to be *poised for the gradient* if the system in Lemma 2 is invertible and well-conditioned.[3] In that case, the factor $\max(h^{-1}|w_{ij}^{\nabla}|)$ can be dropped in equation 2. In the sequel we will always take our meshes poised for both the Laplacian and the gradient.

The approximate normal is then defined as

$$\tilde{n}_i = \frac{(\nabla_i f, -1)^t}{\sqrt{|\nabla_i f|^2 + 1}},$$

and as a consequence of Lemma 2, $\|n_i - \tilde{n}_i\| = O(h^q)$.

3 Discrete Laplace-Beltrami on manifolds

3.1 Normal extension

All we need to know about the Laplace-Beltrami operator on a smooth manifold is the following result.

Lemma 3. *Let f be a scalar C^2 function on a smooth surface $\mathcal{M}$ embedded in $\mathbb{R}^3$, and let $\overline{f}$ be the normal extension of f outside of $\mathcal{M}$, i.e., the extension defined in a thin shell outside of $\mathcal{M}$ by keeping $\overline{f}$ constant along each normals to $\mathcal{M}$. Then, at points in the interior of $\mathcal{M}$,*

$$\Delta_{\mathcal{M}} f = \Delta_{\mathbb{R}^3} \overline{f}. \tag{3}$$

Proof. The distributional definition of the Laplace-Beltrami operator in the interior of $\mathcal{M}$ is

$$\int_{\mathcal{M}} \Delta_{\mathcal{M}} f \, \phi = -\int_{\mathcal{M}} \nabla_{\mathcal{M}} f \cdot \nabla_{\mathcal{M}} \phi,$$

where ∇_M is the gradient intrinsic to $\mathcal{M}$ and ϕ any smooth, compactly supported test function on $\mathcal{M}$. Let ψ be a signed distance function to $\mathcal{M}$, or "level-set" function, obeying locally (near supp ϕ)

$$\psi = 0 \text{ on } \mathcal{M}, \qquad |\nabla \psi| = 1 \text{ outside of } \mathcal{M}. \tag{4}$$

The normal extension $\overline{f}$ of f is then defined locally as

$$\overline{f} = f \text{ on } \mathcal{M}, \qquad \nabla \overline{f} \cdot \nabla \psi = 0 \text{ outside of } \mathcal{M}. \tag{5}$$

It is always possible to solve equations (4) and (5) in a thin shell around the surface, of width $\kappa/2$ where $\kappa = \max_{x \in \mathcal{M}}(\frac{\kappa_1 + \kappa_2}{2})$ is the maximum mean curvature of the surface.

The intrinsic gradient of f is directly related to the gradient in R^3 by $\nabla_{\mathcal{M}} = P\nabla$, where P is projection onto the local tangent plane, $P = I - \nabla\psi \otimes \nabla\psi$. This is true regardless of the way f is smoothly extended outside of $\mathcal{M}$. For the normal extension, this of course becomes

$$\nabla_{\mathcal{M}} f = \nabla \overline{f}.$$

We can now proceed to evaluate $\int_{\mathcal{M}} \nabla_{\mathcal{M}} f \cdot \nabla_{\mathcal{M}} \phi$ by extending f and ϕ normally, and expressing the integral in the embedding space. Let ρ_ϵ be a compactly supported approximate identity in $\mathbb{R}$, i.e., so that ρ_ϵ tends weakly to a Dirac at zero as $\epsilon \to 0$. Assume that the support of ρ_ϵ does not extend outside of a small interval of size $\kappa/2$, uniformly in ϵ. Then $\sigma_\epsilon = \rho_\epsilon \circ \psi$ is an approximate surface measure for $\mathcal{M}$. It follows that

$$\int_M \nabla_{\mathcal{M}} f \cdot \nabla_{\mathcal{M}} \phi = \lim_{\epsilon \to 0} \int_{\mathbb{R}^3} \nabla \overline{f} \cdot \nabla \overline{\phi} \, \sigma_\epsilon \, dx.$$

For fixed ϵ, the right-hand side can be integrated by parts and gives

$$-\int_{\mathbb{R}^3} \Delta \overline{f} \, \overline{\phi} \, \sigma_\epsilon \, dx. \tag{6}$$

[3] In this context, the condition for a neighborhood $\{x_j; j \sim i\}$ to be poised would be that the x_j do not lie on the zero level set of a polynomial of order $p+1$ *with no constant term.*

We have used the fact that $\nabla\sigma_\epsilon = \rho'_\epsilon \nabla\psi$ is perpendicular to $\nabla\overline{f}$, as in equation (5), to get rid of the term involving $\nabla\sigma_\epsilon$. What remains, (6), tends to

$$-\int_{\mathcal{M}} \Delta\overline{f}|_{\mathcal{M}}\, \phi,$$

valid for all smooth ϕ, which establishes the claim. □

Since the extension $\overline{f}$ is normal, all its derivatives along the normal vanish, which allows to reduce $\Delta_{\mathbb{R}^3}$ to differentiation in the tangent plane $T_x\mathcal{M}$:

$$\Delta_{\mathcal{M}} f(x) = \Delta_{T_x\mathcal{M}} \overline{f}(x).$$

Therefore, a discretization of the Euclidean Laplacian in the tangent plane at a point x_i is also a discretization of the Laplace-Beltrami operator at the same point x_i—provided the function $\overline{f}$ on the tangent plane is the normal extension of f on the manifold.

3.2 Main algorithm

Let us denote by n_k the normal vector to $\mathcal{M}$ at points $x_k \in \mathcal{M}$, and, like before, $e_{ik} = x_k - x_i$. Consider the tangent plane $T_{x_i}\mathcal{M}$ to $\mathcal{M}$ at some given point x_i, and the image x'_k of x_k in $T_{x_i}\mathcal{M}$ obtained by following n_k. A bit of algebra shows that the normal projection of x_k is

$$\overline{x}_k = x_k + n_k \frac{e_{ik} \cdot n_i}{n_k \cdot n_i},$$

and since the projection of x_i is itself,

$$\overline{e}_{ik} = e_{ik} + n_k \frac{e_{ik} \cdot n_i}{n_k \cdot n_i}.$$

We are now ready to state the algorithm for the discretized Laplace-Beltrami operator $\tilde{\Delta}_i f$ at a point $x_i \in \mathcal{M}$.

1. Consider a discretization Δ_i of the 2-D Euclidean laplacian. Its weights w_{ij}^{Δ} depend only on the vectors e_{ik}, as we saw in Lemma 1.
2. For each vertex x_k neighboring x_i, compute the normal $\tilde{n}_k$ to $\mathcal{M}$ by the algorithm suggested in Lemma 2.
3. Compute the corresponding new $\tilde{e}_{ik}$ as

$$\tilde{e}_{ik} = e_{ik} + \tilde{n}_k \frac{e_{ik} \cdot \tilde{n}_i}{\tilde{n}_k \cdot \tilde{n}_i}.$$

4. Compute the weigths w_{ij}^{Δ} with $\tilde{e}_{ik}$ in place of e_{ik}.

3.3 Error estimate

As we saw, the normals are known up to some pointwise error of order q. Similarly, the Laplacian in the tangent plane is only known up to an error of order p. Let us see how these errors combine.

Theorem 1. *Fix p and q two positive integers. Consider a sequence of meshes $\{x_i\}$ on $\mathcal{M}$, with diameters $h \to 0$, such that the projections of the respective neighborhoods $\{x_j : j \sim i\}$ on the respective tangent planes are, together, poised for the Laplacian at order p and the gradient at order q. Suppose the Laplacian is computed to order p in the tangent planes, and suppose the computation of the normals is of order q. Then, pointwise and for sufficiently smooth f,*

$$\sup_i |\tilde{\Delta}_i f - \Delta_{\mathcal{M}} f(x_i)| \leq C(f, p, q) \cdot (h^p + h^q).$$

Proof. The approximate computation of the n_k has for first consequence that the 2D Euclidean Laplacian is computed in a wrong plane, going through x_i and slightly tilted with respect to the true tangent plane $T_{x_i}\mathcal{M}$. In $\mathbb{R}^3$ choose coordinates (y_1, y_2, z) so that $T_{x_i}\mathcal{M} \equiv \{z = 0\}$, and (y'_1, y'_2, z') so that the 'wrong plane' is $\{z' = 0\}$. Elementary Taylor expansions of trigonometric functions show that the sets of coordinates can be defined to obey

$$\begin{pmatrix} y'_1 \\ y'_2 \\ z' \end{pmatrix} = \begin{pmatrix} 1 + O(h^{2q}) & O(h^q) & O(h^q) \\ O(h^q) & 1 + O(h^{2q}) & O(h^q) \\ O(h^q) & O(h^q) & 1 + O(h^{2q}) \end{pmatrix} \begin{pmatrix} y_1 \\ y_2 \\ z \end{pmatrix}.$$

The chain rule then gives

$$\frac{\partial^2}{\partial {y'_k}^2} = (1 + O(h^{2q})) \cdot \frac{\partial^2}{\partial y_k^2} + O(h^{2q}) \cdot (\text{other derivatives}), \qquad k = 1, 2.$$

As a result,

$$\sup_i |\Delta_{\mathcal{M}} f(x_i) - \left(\frac{\partial^2}{\partial {y'_1}^2} + \frac{\partial^2}{\partial {y'_2}^2} \right) f(x_i)| \leq C(f, q) \cdot h^{2q}. \tag{7}$$

The second source of error is discretization at order p. As per Lemma 1,

$$\sup_i | \left(\frac{\partial^2}{\partial {y'_1}^2} + \frac{\partial^2}{\partial {y'_2}^2} \right) f(x_i) - \Delta'_i f| \leq C(f, p) \cdot h^p. \tag{8}$$

where Δ'_i is the discrete Laplacian in the tilted plane $z' = 0$.

The last source of error comes from the fact that the other normals $\tilde{n}_k$, for $k \neq i$, also come with an error h^q. Since the diameter of a neighborhood $\{x_j : j \sim i\}$ is bounded by $C(q) \cdot h$, the distance between a point and its normal projection onto the tangent plane obeys $\|x_j - \overline{x}_j\| \leq C(q) \cdot h^2$. Therefore the computed projections $\tilde{x}_j$ obey $\|\tilde{x}_j - \overline{x}_j\| \leq C(q) \cdot h^{2+q}$. It is easier to make this discrepancy bear on the location of evaluation points for the function f, rather than on the computation of the weights w_{ij}^{Δ} through $\overline{e}_{ik}$. So we have

$$|\overline{f}(\tilde{x}_j) - \overline{f}(\overline{x}_j)| \leq C(f, q) \cdot h^{2+q}, \qquad j \sim i$$

Thanks to Definition 1, the magnitude of w_{ij}^{Δ} is at most $O(h^{-2})$, so

$$\sup_i |\Delta'_i f - \tilde{\Delta}_i f| \leq C(f, p, q) \cdot h^q. \tag{9}$$

We conclude by combining (7), (8), and (9). □

Notice that a simple difference formula with $p = q = 2$ is already more accurate than standard off-the-shelf methods like the cotangent formula [12].

4 Discussion

Two practical aspects have been overlooked in the discussion, but may play important roles in applications:

- **Conditioning.** The conditions under which the linear systems in Lemmas 1 and 2 could be ill-conditioned are linked to the geometric arrangement of the points x_j in the neighborhood of each x_i. For instance, they cannot be collinear, and more generally, they should not accumulate near an algebraic curve. I am unaware of any practical, a priori procedure to test for this geometric condition.
- **Positive-definiteness**. The discrete Laplacian need not only be accurate, but also needs to be positive definite. This property is important to avoid instabilities in simulating time-dependent equations involving the Laplacian, like the heat equation. A perturbation argument involving the fact that the eigenvalue of the Laplacian nearest to the origin is of order h^2, whereas numerical errors are of order h^p, $p > 2$, may help settle this question for small h, although in a not-too-elegant manner.

References

[1] M. Belkin, P. Niyogi, Laplacian eigenmaps for dimensionality reduction and data representation *Neural Computation* 15 (2003) 1373–1396

[2] M. Belkin, P. Niyogi, Towards a theoretical foundation for Laplacian-based manifold methods in *Proc. Conf. on Learning Theory*, LNAI 3559 (2005) 486–500

[3] F. Cazals and M. Pouget, Estimating Differential Quantities using Polynomial fitting of Osculating Jets *Computer Aided Geometric Design* 22 (2005) 121–146

[4] D. Cohen-Steiner and J.-M. Morvan, Restricted Delaunay triangulations and normal cycle in *Proc. 19th Annu. ACM Sympos. Comput. Geom.* (2003) 237–246

[5] M. Hein, J.-Y. Audibert, and U. von Luxburg, From graphs to manifolds–Weak and strong pointwise consistency of graph Laplacians in *Proc. Conf. on Learning Theory*, LNAI 3559 (2005) 470–485

[6] J. E. Castillo, J. M. Hyman, M. J. Shashkov, and S. Steinberg, High-order mimetic finite difference method on nonuniform grids in *Proc. ICOSAHOM* (1995)

[7] R. R. Coifman and S. Lafon, Diffusion maps *Appl. Comp. Harm. Anal.* 21:1 (2006) 5–30

[8] S. Lafon, Diffusion maps and geometric harmonics Ph.D. Thesis, Yale University, 2004.

[9] R. R. Coifman and M. Maggioni, Diffusion wavelets *Appl. Comp. Harm. Anal.* 21:1 (2006) 53–94

[10] R. R. Coifman and M. Maggioni, Multiscale Data Analysis with Diffusion Wavelets Technical report YALE/DCS/TR1335, Yale University (2005)

[11] G. Huiskamp, Difference formulas for the surface Laplacian on a triangulated surface *J. Comput. Phys.* 95 (1991) 477–496

[12] M. Meyer, M. Desbrun, P. Schröder, A. H. Barr, Discrete differential-geometry operators for triangulated 2-manifolds in *Proc. Vis. Math III* (2003)

[13] P. J. Olver, On multivariate interpolation *Studies in Applied Mathematics* 116:2 (2006) 201

[14] A. D. Szlam, M. Maggioni, and R. R. Coifman, A General Framework for Adaptive Regularization Based on Diffusion Processes On Graphs Technical report YALE/DCS/TR1365, Yale University (2006)

[15] G. Xu, Discrete Laplace–Beltrami operators and their convergence *Computer Aided Geometric Design* 21 (2004) 767–784