



Sparse Approximation of 3D Meshes Using the Spectral Geometry of the Hamiltonian Operator

Yoni Choukroun¹ · Gautam Pai¹ · Ron Kimmel¹

Received: 7 July 2017 / Accepted: 14 May 2018 / Published online: 24 May 2018
© Springer Science+Business Media, LLC, part of Springer Nature 2018

Abstract

The discrete Laplace operator is ubiquitous in spectral shape analysis, since its eigenfunctions are provably optimal in representing smooth functions defined on the surface of the shape. Indeed, subspaces defined by its eigenfunctions have been utilized for shape compression, treating the coordinates as smooth functions defined on the given surface. However, surfaces of shapes in nature often contain geometric structures for which the general smoothness assumption may fail to hold. At the other end, some explicit mesh compression algorithms utilize the order by which vertices that represent the surface are traversed, a property which has been ignored in spectral approaches. Here, we incorporate the order of vertices into an operator that defines a novel spectral domain. We propose a method for representing 3D meshes using the spectral geometry of the Hamiltonian operator, integrated within a sparse approximation framework. We adapt the concept of a potential function from quantum physics and incorporate vertex ordering information into the potential, yielding a novel data-dependent operator. The potential function modifies the spectral geometry of the Laplacian to focus on regions with finer details of the given surface. By sparsely encoding the geometry of the shape using the proposed data-dependent basis, we improve compression performance compared to previous results that use the standard Laplacian basis and spectral graph wavelets.

Keywords Spectral shape analysis · Mesh compression · Hamiltonian

1 Introduction

In most areas that involve representation of discrete virtual surfaces as 3D meshes, there has been an increasing trend in working with higher precision. This has led to the generation of meshes comprised of a large number of elements, for which the processing, visualization and storage of have become a challenge. The task of transmission of these geometric models over communication networks can lead to a large amount of storage space and put a considerable strain on network resources.

The research on mesh compression started in the direction of single-rate compression methods that constructed a compact representation of an input mesh as a whole. However,

the need for transmission of large geometric models led to the development of algorithms based on progressive mesh compression where a mesh can be transmitted and reconstructed progressively, that is, with increasing levels of detail. This enables the rendering of the 3D model in progression where the finer details of the model can be processed in a cumulative fashion.

The information contained in a mesh is generally divided into two categories: the *geometry* information, which is the position of each vertex of the mesh in the 3D Euclidean space and the *connectivity* or *topological* information, which describes the incidence relations between the mesh vertices. Since the geometric information comprises a dominant part of the mesh, most recent algorithms focused on its efficient compression. Complete reviews, summarizing work on mesh compression, can be found in [2,21,27].

1.1 Related Results

Fourier analysis has been used extensively in many signal processing areas. By projecting the data into the frequency domain, one can retrieve a good approximation of the original

✉ Yoni Choukroun
yonic@cs.technion.ac.il

Gautam Pai
paigautam@cs.technion.ac.il

Ron Kimmel
ron@cs.technion.ac.il

¹ Technion-Israel Institute of Technology, Haifa, Israel

signal with only a few Fourier coefficients. A very prominent application is the JPEG [37] image compression method based on the closely related 2D discrete cosine transform.

Karni and Gotsman [14] proposed a generalization of the Fourier basis on discrete graphs in order to compress mesh vertex positions. This is achieved by projecting the coordinate vectors onto the orthonormal basis obtained from the spectral decomposition of the combinatorial Laplacian of the shape. For smooth meshes, most of the signal energy is contained in the low frequencies of the mesh spectrum, and hence, the coefficients corresponding to the lower frequencies are sufficient to build a good approximation of the original mesh. Ben-Chen and Gotsman [4] proved that for certain geometric models, spectral compression is optimal and is equivalent to principal component analysis.

The idea of using a representative basis for approximating a mesh has been explored in detail in multiple prior works. [15] proposed to use a fixed basis independent of the mesh connectivity. [31] proposed a geometry-aware basis that considers both the connectivity and the geometry of the mesh for mesh compression. Mahadevan [22] replaced the manifold harmonics by a diffusion wavelet bases and showed an improvement in compression performance. This was attributed to the rich multi-scale nature of the wavelet basis.

Methods dealing with basis design for meshes have also been explored in other applications. In order to overcome the non-local support of the Fourier eigenfunctions, compressed eigenfunctions have been adapted from mathematical physics to shape analysis [5,25].

Anisotropic phenomena, such as diffusion, have been explored in different image processing applications [38] and have recently demonstrated impressive results in image compression [28]. Anisotropic Laplacians have been introduced in order to improve the performance compared to the classical isotropic Laplace operator. Kovnatsky et al. [18] modified the metric tensor of the manifold to take into account photometric information. Andreux et al. [3] used principal curvatures in order to construct anisotropic tensor for the task of shape segmentation and shape matching.

Keeping the anisotropic information out of the metric leads to lower order of derivatives, a property that yields numerically more robust operators. Hildebrandt et al. [11] designed a new kind of eigenvibrations constructed as the Hessians of surface energies using extrinsic curvatures and deformation energies. Similarly, Iglesias and Kimmel [13] used artificial surface textures on shapes to define elliptic operators that give birth to a new family of diffusion distances. This operator, also called Hamiltonian, was recently explored in [7], which provides a broad theoretical foundation of the operator and its use in different shape analysis tasks, including data-driven potential optimization, compressed manifold modes and shape matching. In this paper,

we extend this work by analyzing the operator for the task of mesh compression and provide a powerful sparse representation framework based on the Hamiltonian operator, that outperforms previous spectral compression methods.

Given a basis, a plain spectral truncation of the signal in that basis is a fairly simple method for representation. Instead, sparse approximation techniques have been proposed and have become very popular. The basic idea is to estimate a given signal as a *sparse* linear combination of a large pool of constituent vectors—called a dictionary. These vectors, or atoms, are selected such that the coefficients of representation are sparse. The rationale is that high-dimensional signals generally possess intrinsic structures that are better represented in a lower-dimensional linear subspace. Although these dictionaries have traditionally been populated with complete orthogonal bases, redundant or overcomplete dictionaries allow greater flexibility in design by better capturing the intrinsic characteristics of a signal.

The main difficulty with sparse algorithms is the availability of a rich representative dictionary. This seems trivial for signals defined over regularly and consistently sampled domains, like images and speech, but it is not straightforward to extend the idea to non-flat domains like meshes of surfaces in 3D or general graphs. The use of redundant representations for mesh representation and compression has started to emerge in [19,29,32,39]. We extend the ideas discussed in these papers by choosing the eigenfunctions of a data-aware operator as the constituents of a large overcomplete dictionary to represent the mesh.

1.2 Contribution

In this paper, we use the spectral geometry of the Hamiltonian operator for approximation of meshes representing surfaces in 3D. The operator is obtained by modifying the Laplacian with a potential function that defines the rate of oscillation of the harmonics on different regions of the surface. We use a simple and efficient construction of the potential function using a vertex ordering scheme. This modulates the Fourier basis of the 3D mesh to focus on crucial regions of the shape having fine geometric structures equivalent to high frequencies in the Fourier domain. We optimize our approximation using a sparse representation framework by building optimized dictionaries in order to *sparingly* encode mesh geometry [6].

In contrast to [7], which suggests a simplistic truncation scheme for the representation of discrete surfaces, this paper provides an elaborate treatise on the design of the Hamiltonian operator to tackle the specific task of spectral mesh compression. We develop a novel multi-resolution dictionary learning mechanism which takes advantage of the adaptable nature of the Hamiltonian to design a context basis. The combination of the spectral geometry of

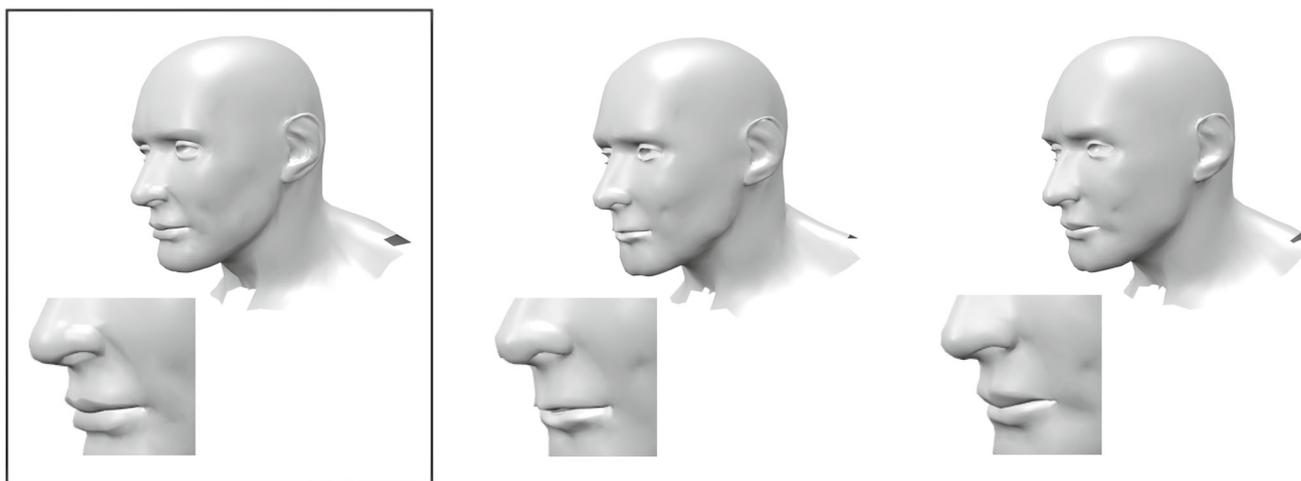


Fig. 1 “The devil lies in the details”. The boxed frame on the left is the original shape. The center and right figures show sparse approximations obtained using the eigenfunctions of the standard Laplacian and the Hamiltonian operators, respectively, using a compression ratio of

1:10. The global Laplacian basis, also called manifold harmonics, fails to capture the structure in the fine details. The approximation using the proposed Hamiltonian basis (figures on the right) preserves this significant fine structure for the same compression ratio

the Hamiltonian in conjunction with a sparse approximation approach outperforms existing *spectral* compression schemes as demonstrated in Sect. 5 and is illustrated in Fig. 1.

2 Background

2.1 Laplace Beltrami Operator

Consider a parameterized manifold $\mathcal{M} : \Omega \subset \mathbb{R}^2 \rightarrow \mathbb{R}^3$ with boundary $\partial\mathcal{M}$. The space of square-integrable functions on \mathcal{M} is denoted by $L^2(\mathcal{M}) = \{f : \mathcal{M} \rightarrow \mathbb{R} \mid \int_{\mathcal{M}} f^2 da < \infty\}$ with the standard inner product $\langle f, g \rangle_{\mathcal{M}} = \int_{\mathcal{M}} fg da$, where da is an area element induced by the Riemannian metric $\langle \cdot, \cdot \rangle_g$.

The Laplace–Beltrami operator (LBO) acting on a scalar function $f \in L^2(\mathcal{M})$ is defined as

$$\Delta_{\mathcal{M}} f \equiv \text{div}_{\mathcal{M}}(\nabla_{\mathcal{M}} f), \tag{1}$$

where $\nabla_{\mathcal{M}}$ represents the intrinsic gradient. The LBO is a self-adjoint operator and hence admits a spectral decomposition: $\{\lambda_i, \phi_i\}$, such that,

$$\begin{aligned} -\Delta_{\mathcal{M}} \phi_i &= \lambda_i \phi_i, \\ \langle \phi_i, \phi_j \rangle_{\mathcal{M}} &= \delta_{ij}, \end{aligned} \tag{2}$$

where $\lambda_i \in \mathbb{R}$, $0 = \lambda_1 \leq \lambda_2 \leq \dots$ and δ_{ij} the Kronecker delta function. In case \mathcal{M} has a boundary, we add Neumann boundary condition to ϕ_i . The LBO eigendecomposition can be obtained by minimizing the Dirichlet energy

$$\begin{aligned} \min_{\phi_i} \int_{\mathcal{M}} \|\nabla_{\mathcal{M}} \phi_i\|_g^2 da, \\ \text{s.t.} \quad \langle \phi_i, \phi_j \rangle_{\mathcal{M}} = \delta_{ij}. \end{aligned} \tag{3}$$

Therefore, the LBO eigenfunctions can be seen as an extension of the *smooth* Fourier harmonics in Euclidean space to manifolds and are also referred to as *manifold harmonics*[20,36].

2.2 Graph Laplacian

Let $\mathcal{G} = (\Pi, \Upsilon)$ be an undirected graph of n vertices $\Pi = \{\pi_i\}_{i=1}^n$ representing a sampled version of a manifold \mathcal{M} . Let the graph edges be equipped with nonnegative weights $\{w_{ij}\}_{i,j=1}^n$, defining the weighted adjacency matrix W of the graph as

$$[W]_{ij} = \begin{cases} w_{ij}, & \text{if } (i, j) \in \Upsilon, \\ 0, & \text{otherwise.} \end{cases} \tag{4}$$

The diagonal vertex degree matrix A is defined as

$$[A]_{ii} = a_{ii} = \sum_j w_{ij}. \tag{5}$$

The unnormalized graph Laplacian is defined as the $n \times n$ matrix $L = A - W$. For properties of the Laplacian, the reader is referred to [8,24].

The combinatorial graph Laplacian is constructed by solely considering the connectivity of the mesh, where the edge weights are defined as $w_{ij} = 1$. Therefore, the degree matrix A is populated with the valence of each vertex. While

other *geometric* discretizations of the Laplacian on graphs exist [30], the combinatorial Laplacian is used in this paper, since the mesh topology can be efficiently encoded .

The discrete manifold harmonics are then obtained by solving the spectral decomposition of the combinatorial Laplacian

$$L\Phi = \Phi\Lambda. \tag{6}$$

Here, $\Phi \in \mathbb{R}^{n \times n}$ is the matrix of eigenvectors where each column ϕ_k is one eigenvector and Λ is the diagonal matrix of eigenvalues where each element λ_k is one eigenvalue.

2.3 Spectral Mesh Compression

Given a complete basis $\{\varphi_i\}_{i=0}^\infty$ on \mathcal{M} , any function $f \in L^2(\mathcal{M})$ can be expressed as

$$f = \sum_{i=1}^\infty \langle f, \varphi_i \rangle_{\mathcal{M}} \varphi_i = \sum_{i=1}^\infty \hat{f}_i \varphi_i, \tag{7}$$

where \hat{f}_i is the i th (manifold) Fourier coefficient of f . Since the manifold harmonics form a complete basis, Karni and Gotsman [14] proposed to use them for spectral mesh compression. Considering the mesh coordinate functions X, Y, Z as functions defined on the vertices of the mesh, the basic idea of spectral mesh compression is to compute the Fourier transform of the coordinate functions and truncate the high-frequency coefficients. Thus, if u is a coordinate function, it can be approximated using the graph Laplacian decomposition as

$$u \approx \sum_{i=1}^K \langle u, \phi_i \rangle_{\mathcal{G}} \phi_i \tag{8}$$

with $K \ll n$ and $\langle u, \phi_i \rangle_{\mathcal{G}} = \phi_i^T u$. By using the graph connectivity information and constructing the manifold harmonics as a representative basis, the geometric information of the mesh can be encoded efficiently through the obtained coefficients. Figure 2 illustrates spectral reconstructions as applied to a toy example of a 2D curve.

Simple truncation as shown in Eq. (8) uses a restrictive assumption of only focusing on the lower frequencies and is analogous to low-pass filtering of the signal. In general, the lower frequencies encapsulate most of signal energy, and the reconstructed mesh tends to preserve the global features of the original. However, local geometry and fine details of the mesh corresponding to high frequencies are generally missing and require a high number of eigenvectors for their preservation if a truncation approach is followed.

In the next section, we propose a more pragmatic approach of sparsity where the data itself searches for its best con-

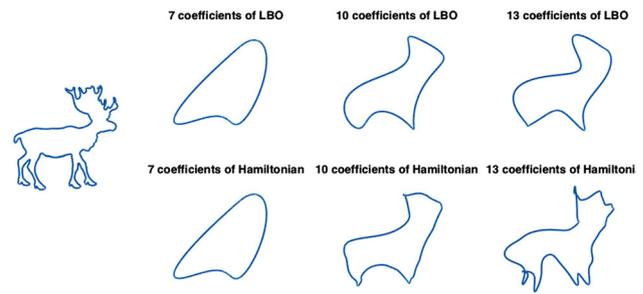


Fig. 2 Approximation of a planar curve by simple truncation of the LBO and the proposed basis (Hamiltonian). The potential function for the Hamiltonian is defined by the error between the curve and its LBO truncated approximation. By biasing the 1D harmonics with this potential, the reconstruction with the Hamiltonian can be seen to preserve finer details of the contour for the same number of coefficients

stituents (frequencies) with the objective of attaining sparsity in a certain basis of representation. This leads to an improved compression strategy.

2.4 Sparse Approximation of Mesh Coordinates

As motivated in the last section, classical methods are based on the premise that family of signals are best represented using only the first few components of an optimal orthonormal basis. The concept of sparsity provides an alternative perspective for representation.

Consider a graph \mathcal{G} and $D \in \mathbb{R}^{n \times m}$ a normalized over-complete ($m \gg n$) dictionary containing m atoms $d_i \in \mathbb{R}^n$. We aim to approximate any given signal u on \mathcal{G} by

$$u = D\alpha = \sum_{i=1}^m d_i \alpha_i, \tag{9}$$

with $\alpha = [\alpha_1, \dots, \alpha_m]^T$ the coefficient representation of the input signal u with respect to the dictionary D . In order to achieve significant reduction in storage, we should assume that the number of nonzero elements of the coefficient vector α should satisfy $\|\alpha\|_0 = k \ll n$. Thus, the sparse approximation of the signal can be obtained by solving the sparse coding problem

$$\begin{aligned} \min_{\alpha} & \|u - D\alpha\|_2^2 \\ \text{s.t.} & \|\alpha\|_0 = k. \end{aligned} \tag{10}$$

Let us be given a multiple channel signals $U \in \mathbb{R}^{n \times c}$, the new coefficient matrix is designated by $\Gamma = [\alpha_1, \dots, \alpha_c]^T$ with $U = D\Gamma$. The new sparse approximation is obtained by solving the simultaneous sparse approximation

$$\begin{aligned} \min_{\Gamma} & \|U - D\Gamma\|_F^2 \\ \text{s.t.} & \|\alpha_1\|_0 = \dots = \|\alpha_c\|_0 = k, \end{aligned} \tag{11}$$

where $\|\cdot\|_F$ denotes the Frobenius norm. In our mesh compression configuration, the input signal is defined by the mesh coordinates such that $U = [X, Y, Z]$.

Since sparse coding solution is a combinatorial NP-hard problem, well-known relaxations are generally used [23,26]. In this paper, Simultaneous Orthogonal Matching Pursuit algorithm [33] is used for sparse approximation of the mesh geometry. Pursuit algorithms are based on a greedy approach of solving for the optimization in (11) by incrementally selecting atoms from the dictionary D that minimize the residual error. The orthogonal matching pursuit is one such way of picking the atoms where at each stage of the selection, the residual error vector is orthogonal to the atoms already selected in the support and hence will not be chosen again. Simultaneous OMP is a scheme of atom selection such that multiple signals *have the same support* in the dictionary D , i.e., simultaneously solving three sparse problems by imposing same support constraint.

The success of a sparsity-based algorithm depends on the availability of a representative dictionary. However, developing dictionaries for the representation of functions defined over graphs is a non-trivial task, since they do not possess a fixed structure as in Euclidean domains (grid). As a result, elements composed from fixed spectral bases have been used for populating the atoms of the dictionary [19].

The manifold harmonics can be used because of their stability and compactness, that is, low frequencies contain most of the signal energy. However, fine local structures are not well represented using this basis due to the global nature of their support. In [39], it was shown that a dictionary constructed from spectral graph wavelets [10] shows better compression performance as compared to the standard Laplacian basis, since it encapsulates fine details and some dominant high-frequency structures of the mesh. However, since the locality of the basis cannot be controlled, the wavelet dictionaries are highly overcomplete, yielding large dictionaries which enforce higher compression rates. These shortcomings motivate the need for finding a basis which can be modulated in a data-dependent manner. This will lead to an optimal local–global structure trade-off, thereby requiring a smaller cardinality of representation. In the next section, we borrow an operator from quantum mechanics which has precisely the aforementioned properties.

3 Hamiltonian Operator

3.1 Definition

A Hamiltonian operator H , also called Schrödinger operator, is an operator acting on a scalar function $f \in L^2(\mathcal{M})$ on a manifold \mathcal{M} that has the form

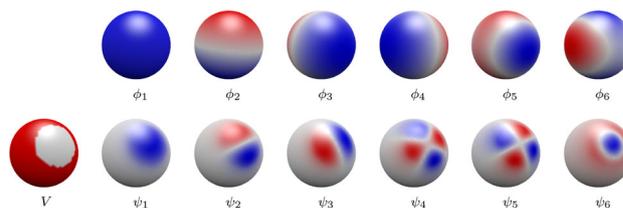


Fig. 3 First eigenfunctions of the LBO on a sphere (top). Potential function defined on the sphere and the corresponding Hamiltonian basis (bottom). Hot and cold colors represent positive and negative values, respectively. The Hamiltonian basis concentrates the low harmonics to the low potential region

$$Hf = -\Delta_{\mathcal{M}}f + \mu Vf, \tag{12}$$

where $V : \mathcal{M} \rightarrow \mathbb{R}$ is a potential function and $\mu \in \mathbb{R}$. The Hamiltonian plays a fundamental role in the field of quantum mechanics via the famous Schrödinger equation that describes the wave motion of a particle. The spectral decomposition $\{\psi_i, E_i\}_{i=0}^{\infty}$ of the Hamiltonian is defined by

$$H\psi_i = E_i\psi_i \tag{13}$$

where E_i denotes the energy of a particle at the stationary eigenstate ψ_i . Here, $\psi_i(x)$ represents the wave function of the particle such that $|\psi_i(x)|^2$ is interpreted as the probability distribution of finding the particle at a given position x . Since the Hamiltonian is a symmetric operator, its eigenfunctions form a complete orthonormal basis on the manifold \mathcal{M} . An illustration of the Hamiltonian basis is given in Fig. 3.

The Hamiltonian eigendecomposition from (13) is obtained from the Euler–Lagrange equation of

$$\begin{aligned} \min_{\psi_i} \quad & \int_{\mathcal{M}} \left(\|\nabla_{\mathcal{M}}\psi_i\|_g^2 + \mu V\psi_i^2 \right) da, \\ \text{s.t.} \quad & \langle \psi_i, \psi_j \rangle_{\mathcal{M}} = \delta_{ij}. \end{aligned} \tag{14}$$

The parameter μ controls the potential energy and defines the trade-off between local and global support of the basis. Larger values of μ will give solutions that concentrate on the low potential regions, while a smaller μ will give solutions that better minimize the total energy at the expense of more extended wave functions.

In its discrete setting, the Hamiltonian basis is obtained by solving

$$H\psi_i = (L + \mu V)\psi_i = E_i\psi_i \tag{15}$$

with $\{\psi_i\}_{i=1}^n \in \mathbb{R}^n$. Here L denotes the previously described graph Laplacian, and V a diagonal matrix that is defined by the potential scalar function values at vertices $\pi_i \in \Pi$. Since V only modifies the diagonal of the sparse Laplacian, the Hamiltonian is also a sparse matrix with the same nonzero

entries. Thus, there is no increase in the computational cost of the decomposition compared to that of the Laplacian.

3.2 Analysis and Visualization of the Operator

Among the numerous reasons that motivated the selection of the Laplacian basis for shape analysis, one ubiquitous is its efficiency in representing smooth functions. This property is defined by the following observation for functions with L^2 bounded gradient magnitudes:

$$\|r_n\|_2^2 = \left\| f - \sum_{i=1}^n \langle f, \phi_i \rangle \phi_i \right\|_2^2 \leq \frac{\|\nabla_g f\|_g^2}{\lambda_{n+1}}. \tag{16}$$

This result can be established as follows,

$$\begin{aligned} \|\nabla_g f\|_g^2 &= \int_{\mathcal{M}} (-\Delta_{\mathcal{M}} f) f \, da \\ &= \sum_{i=1}^{\infty} \int_{\mathcal{M}} (\langle f, \phi_i \rangle_{\mathcal{M}} \lambda_i \phi_i) f \, da = \sum_{i=1}^{\infty} \lambda_i \langle f, \phi_i \rangle_{\mathcal{M}}^2 \\ &\geq \sum_{i=n+1}^{\infty} \lambda_i \langle f, \phi_i \rangle_{\mathcal{M}}^2 \geq \lambda_{n+1} \\ &\quad \sum_{i=n+1}^{\infty} \langle f, \phi_i \rangle_{\mathcal{M}}^2 = \lambda_{n+1} \|r_n\|_2^2. \end{aligned} \tag{17}$$

The representation error r_n for a function f using the Laplacian basis $\{\phi_i\}_{i=1}^{\infty}$ is bounded by a factor of the squared magnitude of its gradient. This result was subsequently proved to be optimal for representing smooth functions over surfaces in [1], which says that there exists no other basis with such a bound for all possible $L^2(\mathcal{M})$ bounded gradient functions.

To gain further insight, we simulate this result in a discrete setting by analyzing the gradient operators for a simple one-dimensional case. In matrix form, since L is a positive semi-definite matrix, there exists a matrix D such that $L = D^T D$ and the discrete equivalent of result 17 is:

$$\|r_n\|_2^2 = \left\| f - \sum_{i=1}^n \langle f, \phi_i \rangle \phi_i \right\|_2^2 \leq \frac{\|Df\|_2^2}{\lambda_{n+1}}, \tag{18}$$

It is interesting to perform a similar analysis on the eigenfunctions of the Hamiltonian, and Fig. 4 shows a visualization of the gradient operator of the Laplacian and the induced gradient of the Hamiltonian. The result in Eq. (18) gets modified to account for the action of the potential.

$$\|r_n\|_2^2 = \left\| f - \sum_{i=1}^n \langle f, \psi_i \rangle \psi_i \right\|_2^2 \leq \frac{\|W Df\|_2^2}{E_{n+1}}, \tag{19}$$

with $W = (I + D^{-T} V D^{-1})^{\frac{1}{2}}$ such that $H = (W D)^T W D$. Thus, the Hamiltonian basis can be considered as optimal

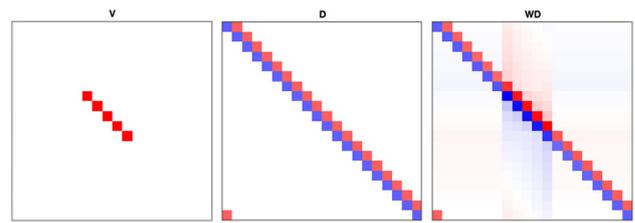


Fig. 4 Matrix visualizations of the potential function (left), the standard gradient matrix D (middle) and the weighted gradient matrix $WD = (I + D^{-T} V D^{-1})^{\frac{1}{2}} D$ (right) for Euclidean 1D example

for the representation of functions which have a bounded *weighted* gradient. The weights depend on the potential function and its action in different regions.

The gradient matrix D for the 1D case is the result of a simple finite difference scheme of one-dimensional forward differentiation yielding a circulant populated with the kernel $[1, -1]$ circulated along the diagonal. The Hamiltonian has the effect of escalating the cost of a strong gradient *exclusively in regions of high potential*. This shows that coefficients of the induced gradient matrix are larger in the high potential regions as compared to anywhere else in the domain. The nonzero values of the off-diagonal elements suggest that the linear operation loses its shift invariance property and the effect of a derivative is no more an exclusive property of its local neighborhood but also depends on its global positioning imposed by the potential function. Thus, the Hamiltonian operator advocates measuring smoothness differently for different regions of the domain, and this is a useful property to exploit in a compression setup.

The Hamiltonian operator can improve the compression performance by modifying the harmonics in order to emphasize *designated* regions of interest. We propose a method for choosing a potential function which does not require any additional encoding. We construct a dictionary based on the eigendecomposition of both the Laplacian and the Hamiltonian that can focus on difficult reconstruction areas of the harmonics. Thus, by designing high vibrations in selected regions of the shape, our dictionary is much less redundant than the wavelets proposed in [22], has a better ability to encode high-frequency details and achieves better compression performance.

4 Hamiltonian Operator for Spectral Mesh Compression

4.1 Potential Design

Consider a mesh (graph) $M = \{\Pi, \Upsilon\}$ with $|\Pi| = n$ and the coordinates matrix U . Given the Laplacian eigenvectors matrix $\Phi \in \mathbb{R}^{n \times n}$, one can design a fixed dictionary D_L

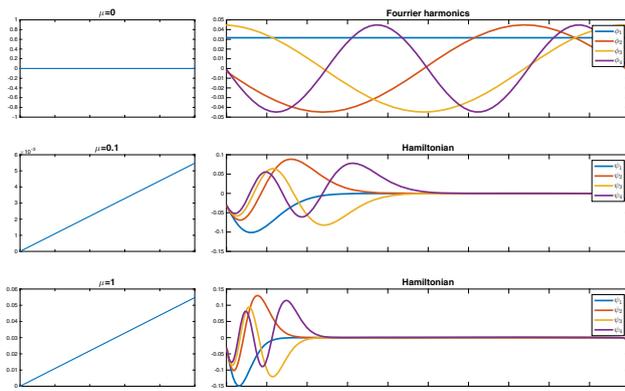


Fig. 5 Zero potential (left up) and its corresponding harmonics (right up). Linear potential for different parameter μ (left middle and bottom) and their corresponding first eigenfunctions (right middle and bottom). The Hamiltonian eigenfunctions focus on the low potential areas without neglecting other regions for higher energies

composed of the full harmonics basis matrix $D_L = \Phi$ and solve the simultaneous sparse coding problem of (11) where $\tilde{U} = D_L \Gamma \approx U$. Then, a L_2 norm reconstruction error $\epsilon(\pi_i) = \|\pi_i - \hat{\pi}_i\|_2$ can be obtained at each vertex of the mesh. This vertex error profile over the mesh provides an indication of difficult regions that can be enhanced through the potential.

Since a fixed potential must be used to avoid additional encoding, we permute the arbitrary indexing of the mesh according to this error. The vertices are sorted in ascending order such that vertices having small indexes correspond to points with a small error, while vertices with the larger index correspond to regions with higher distortion in approximation. It can be done trivially in $\mathcal{O}(\mathcal{Y})$ complexity and has no influence on the mesh itself since the ordering has no physical meaning.

Given the Hamiltonian $H = L + \mu V$, we propose to use a fixed potential matrix defined as

$$V = \text{diag}(1, \dots, n), \tag{20}$$

normalized such that $\|V\|_F^2 = \sum_i^n V_{ii}^2 = 1$. We illustrate the influence of a linear potential in one dimension in Fig. 5. Since we sorted the indexes according to the error, a designed hierarchical energy of the potential is defined over the shape, while the potential encoding itself is avoided. At the decoder side, the Hamiltonian can be built using the mesh connectivity (Laplacian) and the encoded scalar μ coupled with the fixed linear potential V described above.

4.2 Dictionary Construction

We construct a dictionary built from the eigendecomposition of both the Laplacian and the Hamiltonian to benefit from their global and local properties. The potential is defined

according to the representation error of the Laplacian basis as described in the previous section. We encode multiple constants μ in order to obtain a multi-resolution of the basis. The regularization constants are found via a direct search on a given domain.

Thus, our dictionary can be obtained as:

$$D_S = [\Phi, \Psi_{\mu_1}, \dots, \Psi_{\mu_S}], \tag{21}$$

where the sub-dictionary $\Psi_{\mu_i} \in \mathbb{R}^{n \times n}$ is the matrix of eigenvectors of the Hamiltonian $L + \mu_i V$ and $\Phi = \Psi_0$. The constant μ_i is obtained by solving

$$\begin{aligned} \min_{\mu_i} \quad & \min_{\Gamma} \|U - D_i \Gamma\|_F^2 \\ \text{s.t.} \quad & \|\alpha_1\|_0 = \dots = \|\alpha_c\|_0 = k, \end{aligned} \tag{22}$$

The algorithm is summarized in Algorithm 1. Given a compression ratio requirement, the number of atoms is calculated according to the formulas elucidated in Sect. 5.2.

Algorithm 1: Algorithm for sparse approximation

Input : Mesh connectivity \mathcal{Y} , mesh coordinates U and compression ratio k

Output: Sparse coding $\{\alpha_i\}_i$ and coefficients $\{\mu_j\}_j$

- 1 Compute the pointwise error $\epsilon(\pi_i) = \|\tilde{U}(\pi_i) - U(\pi_i)\|_2$ where $\tilde{U} = \Phi \Gamma$ is obtained from (11);
 - 2 Permute the vertices according to the sorting of the error ϵ ;
 - 3 Build constant potential V as in (20) and set $j = 0$;
 - 4 **do**
 - 5 $j = j + 1$;
 - 6 Find μ_j minimizing (22) using $D_j = [\Phi, \Psi_{\mu_1}, \dots, \Psi_{\mu_j}]$;
 - 7 **while** Decrease in representation error;
-

4.3 Discussion

Algorithm 1 is designed to take advantage of the two principle components of our contribution: sparsity- and data-dependent basis. The potential design described in Sect. 4.1 alters the Laplacian basis to focus on error-prone regions of the shape. The extent to which such an alteration must be enforced is encoded in the coefficients: $\{\mu_j\}_{j=1}^{j=K}$ which are sequentially optimized in Algorithm 1. The final dictionary is comprised of the regular Laplacian eigenvectors appended with these Hamiltonian eigenvectors. The results in Sect. 5 demonstrate that this provides a distinct improvement in shape-reconstruction errors.

Another advantage of our approach is that the potential design is based on mesh connectivity information alone. The vertex re-ordering enables us to encode information about the error-prone regions of the shape, which are encoded in the eigenvectors of the Hamiltonian. This re-ordering scheme

avoids any secondary encoding of information and hence provides a distinct gain in saving of information units for compression.

5 Experimental Results

5.1 Evaluation Metric

In order to measure the loss resulting from an estimated approximation, Karni and Gotsman [14] used a visual metric that captures the visual difference between the original mesh $M = \{\Pi, \mathcal{Y}\}$ and its reconstruction $\tilde{M} = \{\tilde{\Pi}, \tilde{\mathcal{Y}}\}$. The measure is a linear combination of the RMS geometric distance between corresponding vertices in both models and a visual metric capturing the smoothness of the reconstruction. The error $\epsilon \in \mathbb{R}^n$ at each vertex is defined as

$$\epsilon_M(\tilde{\pi}_i) = \frac{1}{2n} \left(\|\pi_i - \tilde{\pi}_i\|_2 + \|GL(\pi_i) - GL(\tilde{\pi}_i)\|_2 \right). \quad (23)$$

The visual metric GL is defined as

$$GL(\pi_i) = \pi_i - \frac{\sum_{j \in \mathcal{N}(i)} l_{ij}^{-1} \pi_j}{\sum_{j \in \mathcal{N}(i)} l_{ij}^{-1}} \quad (24)$$

with $\mathcal{N}(i)$ is the set of indexes of the neighbors of vertex i , and l_{ij} is the geometric distance between vertices i and j . A global representation error is obtained by summing ϵ_M over the mesh vertices. We normalize the error by the area of the shape for uniform comparison between different shapes.

5.2 Compression Ratio

The compression ratio calculation is obtained as a ratio of the net information before and after the dictionary encoding [39]. If each coordinate takes up to k bits, then for a mesh of n vertices, the total uncompressed information is $3nk$. Now, assuming that the dictionary $D \in \mathbb{R}^{m \times n}$ is known in advance, or can be constructed from the connectivity information of the mesh, on both the encoder and decoder sides, let n_d be the total number of selected atoms each taking k_d bits to encode each coordinate function, making the total information equal to $3n_d k_d$. The information is required to store the indexes of the coefficients which is given by $\min(m, n_d \lceil \log_2 m \rceil)$. The compression ratio is given as

$$\text{C.R.} = \frac{3n_d k_d + \min(m, n_d \lceil \log_2 m \rceil)}{3nk}. \quad (25)$$

The min function is used since it can be more efficient to represent the selected atoms with a bit vector of size m instead

of the indexes themselves. The compression ratio in the coefficient truncation scheme is given by

$$\text{C.R.} = \frac{n_d}{n}, \quad (26)$$

since only the ordered coefficients must be encoded. The compression ratio of the Hamiltonian dictionary can be easily extended by adding $n_\mu k_\mu$ to the numerator of (25), with n_μ the number of regularization coefficients of the Hamiltonian and k_μ the number of bits required to encode each one of them.

5.3 Shape Partition

Spectral methods, including the one proposed in this paper, are hinged on the eigendecomposition of a large matrix. This $\mathcal{O}(n^3)$ time complexity operation can be computationally demanding, especially when the number of points is large. Therefore, as motivated in [14], we resort to mesh partitioning, where we segment a mesh into smaller constituents and compress each segment independently. We use the METIS [16] package for fast graph partitioning and segmentation. Therefore, the time complexity of the proposed framework can be seen as *linear* in the number of partitions of fixed size.

The Neumann boundary conditions induced by the graph Laplacian may cause a discrepancy between the reconstructed segments along shared boundaries. Mesh partitioning may result in a so-called *edge effects* [14], a degradation of the reconstruction along the sub-mesh boundaries. However, compressing smaller meshes separately has the advantage of better capturing local properties of the mesh even with low number of coefficients. Also, from our experiments, this degradation is not visually sensitive even for low compression rate and especially not with the suggested method. We partition the mesh into sub-meshes of few hundred vertices. This ensures both faster computation of the coefficients and fewer difficult areas for reconstruction the Hamiltonian operator can focus on.

5.4 Results

Figures 6 and 7 show a visual comparison of the shape approximation results. We propose two visualizations. The first row presents a regular reconstruction with the three different methods for a fixed compression ratio. The second row shows a profile of the error plot over the original mesh to highlight regions of erroneous representation. Clearly, the MHB has a fairly large error profile over the surface. The results with the spectral graph wavelets show an improvement as compared to the manifold harmonics. However, the sparse approximation with the Hamiltonian basis shows a

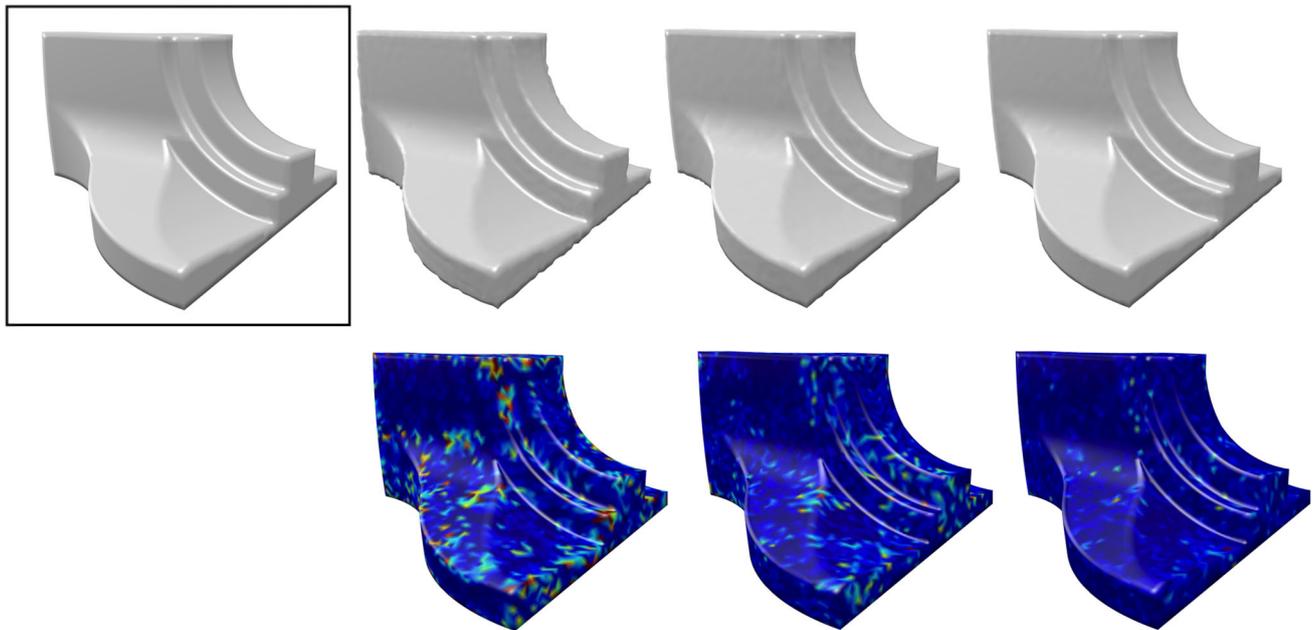


Fig. 6 Fandisk model (boxed mesh) composed of 12,946 vertices. (Top) Shape reconstruction for a compression ratio of 6:10 using dictionaries composed of the MHB, MHB + SGW and MHB + Hamiltonian (from

left to right). (Bottom) Reconstruction error using the proposed spectral bases. Cold and hot colors represent nonnegative low and high values, respectively

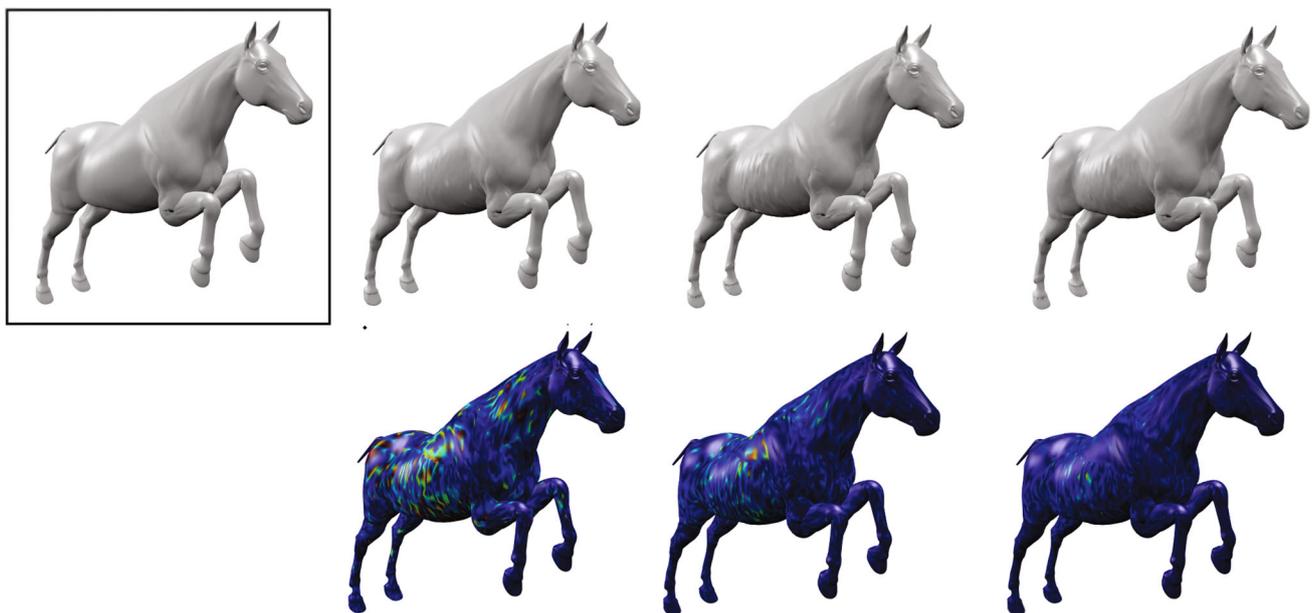


Fig. 7 Horse model (boxed mesh) composed of 19248 vertices. (Top) Shape reconstruction for a compression ratio of 4:10 using dictionaries composed of the MHB, MHB + SGW and MHB + Hamiltonian (from left to right). (Bottom) Reconstruction error using the proposed spectral bases

distinct positive difference as compared to the other two methods both in the reconstruction and in the error profiles.

Figures 8, 9 and 10 show compression results of our method compared with other spectral techniques on different shapes. We present results according to four prominent

spectral methods. First with the plain truncation scheme of the Manifold Harmonic Basis (MHB) as suggested in [14]. Second, the truncation scheme of the Hamiltonian with optimal potential as suggested in [7]. Third, MHB with sparse coding, i.e., using the S-OMP algorithm (MHB-SOMP) on the manifold harmonics dictionary. Fourth, we present Spec-

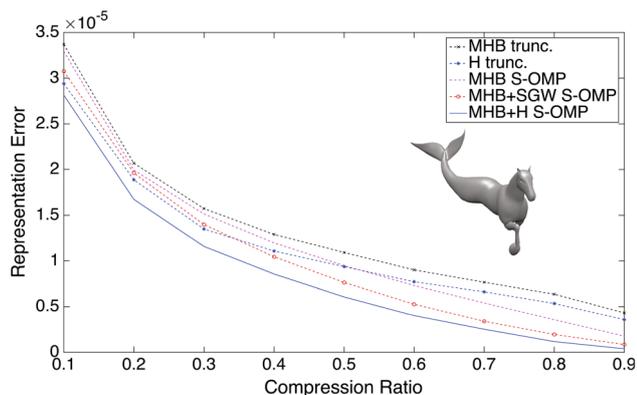


Fig. 8 Reconstruction error as a function of the compression ratio for the seahorse model (2194 vertices)

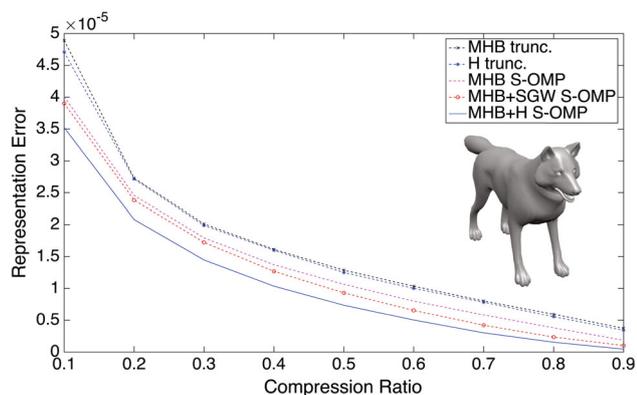


Fig. 9 Reconstruction error as a function of the compression ratio for the wolf model (4344 vertices)

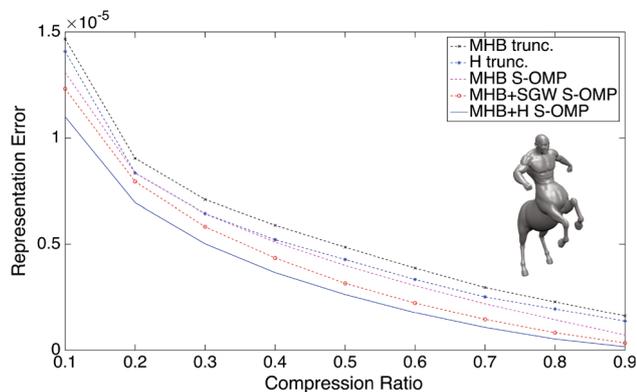


Fig. 10 Reconstruction error as a function of the compression ratio for the centaur model (15768 vertices)

tral graph wavelets with sparse coding (MHB + SGW – SOMP) using dictionaries defined with the harmonics and the graph wavelets as presented in [39]. We used 32 bits as for single bit-precision in the computation of compression ratio. An illustration of the mesh partition result and recon-

struction performance between the methods is presented in Fig. 11.

We can see that our algorithm shows a considerable improvement in compression performance by generating smaller errors over the entire span of compression requirements for different shapes, even under non-uniform triangulations. We highlight the fact we leverage benefits in optimizing on both fronts: a superior basis and the improved sparse coding mechanism for representation. Using only the Hamiltonian basis in a simple truncation scheme as suggested in [7] shows improvement over standard spectral techniques making it comparable to the pursuit-based solutions, but the major gain comes only after its use in conjunction with sparse representation.

Finally, in order to gain a global assessment of the performance of our method in comparison with generic non-spectral mesh compression algorithms (motivated from [21]), we compiled an average bits-per-vertex (bpv) analysis in Table 1, where the proposed method bpv is obtained on the same shapes tested in [14] and with the same reconstruction error threshold. However, since other prominent non-spectral methods like [9,17] involve some form of remeshing, the comparison may not be entirely appropriate, but provides a high-level comparison between the compression types. The general comparison shows that spectral algorithms show lower compression rates as compared to methods like [17] and [9]. However, our approach of using the spectral route along with sparse representations shows considerable gains without any assumption on the mesh connectivity.

6 Conclusion and Future Work

We introduced the Hamiltonian operator for spectral mesh compression. This operator uses the concept of a potential function whose action leads to the modulation of the spectral geometry of that domain. We use this operation to modify the standard manifold harmonics of the 3D surface by emphasizing the role of regions with errors that emerge as a result of compression with classical harmonics. We maximize compression performance by constructing meaningful dictionaries from this basis in a sparse approximation framework and show a distinct improvement from previous spectral techniques.

In future work, we intend to explore the design and construction of newer meaningful potential functions for 3D surfaces. The use of the Schrödinger operator for generic graph structures or discrete manifolds is an uncharted territory that can lead to interesting spectral embeddings.



Fig. 11 Michael model after mesh partition (boxed mesh) composed of 52565 vertices followed by the shape-reconstruction error map for a compression ratio of 1:10 using MHB truncation and dictionaries com-

posed of the MHB + SGW and MHB + Hamiltonian (from left to right). The error scale of the MHB truncation is ten times bigger than the others

Table 1 Comparison of the suggested method with the main progressive mesh compression algorithms

| Algorithm | Total compression rates (bpv) | Compression type |
|----------------------------------------|-------------------------------|----------------------------------|
| Progressive meshes [12] | 37 | Vertex split |
| Wavemesh [35] | 19 | Irregular Wavelets |
| Spectral compression [14] | 19 | Spectral |
| Incremental parametric refinement [34] | 15 | Parametric refinement |
| Spectral Hamiltonian (proposed method) | 12 | Spectral |
| Wavelet compression [17] | 8 | <i>Remeshing based</i> |
| Normal meshes [9] | 6 | <i>Mesh-Sequence compression</i> |

The bold enlighten the performance of the proposed method (best results among related methods (above) and slightly worse than other less restrictive approaches (below))

References

- Aflalo, Y., Brezis, H., Kimmel, R.: On the optimality of shape and data representation in the spectral domain. *SIAM J. Imaging Sci.* **8**(2), 1140–1160 (2015)
- Alliez, P., Gotsman, C.: Recent Advances in Compression of 3D Meshes. In: Dodgson, N.A., Floater, M.S., Sabin, M.A. (eds.) *Advances in Multiresolution for Geometric Modelling*. Springer, Berlin, Heidelberg (2005)
- Andreux, M., Rodola, E., Aubry, M., Cremers, D.: Anisotropic Laplace–Beltrami operators for shape analysis. In: *NORDIA'14-Sixth Workshop on Non-Rigid Shape Analysis and Deformable Image Alignment* (2014)
- Ben-Chen, M., Gotsman, C.: On the optimality of spectral compression of mesh data. *ACM Trans. Graph. (TOG)* **24**(1), 60–80 (2005). <https://doi.org/10.1145/1037957.1037961>
- Bronstein, A., Choukroun, Y., Kimmel, R., Sela, M.: Consistent discretization and minimization of the ℓ_1 norm on manifolds. In: *2016 Fourth International Conference on 3D Vision (3DV)*. IEEE (2016)
- Choukroun, Y., Pai, G., Kimmel, R.: Schrödinger Operator for Sparse Approximation of 3D Meshes. In: Brentzen J.A., Hildebrandt K. (eds.) *Symposium on Geometry Processing 2017-Posters*. The Eurographics Association (2017)
- Choukroun, Y., Shtern, A., Bronstein, A., Kimmel, R.: Hamiltonian operator for spectral shape analysis. [arXiv:1611.01990](https://arxiv.org/abs/1611.01990) (2016)
- Chung, F.R.K.: *Spectral graph theory*: [CBMS conference on recent advances in spectral graph theory held at California State University at Fresno, June 6–10, 1994]. CBMS Regional Conference Series in Mathematics. American Mathematical Society, Providence (1997)
- Guskov, I., Vidimčec, K., Sweldens, W., Schröder, P.: Normal meshes. In: *Proceedings of the 27th Annual Conference on Computer Graphics and Interactive Techniques*. ACM Press/Addison-Wesley Publishing Co., New York (2000)
- Hammond, D.K., Vanderghenst, P., Gribonval, R.: Wavelets on graphs via spectral graph theory. *Appl. Comput. Harmon. Anal.* **30**, 129–150 (2011)
- Hildebrandt, K., Schulz, C., von Tycowicz, C., Polthier, K.: Modal shape analysis beyond Laplacian. *Comput. Aided Geom. Des.* **29**, 204–218 (2012)
- Hoppe, H.: Progressive meshes. In: *Proceedings of the 23rd Annual Conference on Computer Graphics and Interactive Techniques*. ACM (1996)
- Iglesias, J.A., Kimmel, R.: Schrödinger Diffusion for Shape Analysis with Texture. In: Fusiello, A., Murino, V., Cucchiara, R. (eds.) *Computer Vision – ECCV 2012. Workshops and Demonstrations: Florence, Italy, October 7–13, 2012, Proceedings, Part I*. Springer, Berlin, Heidelberg (2012)
- Karni, Z., Gotsman, C.: Spectral compression of mesh geometry. In: *Proceedings of the 27th Annual Conference on Computer Graphics and Interactive Techniques*. ACM Press/Addison-Wesley Publishing Co., New York (2000)

15. Karni, Z., Gotsman, C.: 3d mesh compression using fixed spectral bases. In: Proceedings of Graphics Interface 2001. Canadian Information Processing Society, Mississauga (2001)
16. Karypis, G., Kumar, V.: A fast and high quality multilevel scheme for partitioning irregular graphs. *SIAM J. Sci. Comput.* **20**, 359–392 (1998)
17. Khodakovsky, A., Schröder, P., Sweldens, W.: Progressive geometry compression. In: Proceedings of the 27th Annual Conference on Computer Graphics and Interactive Techniques. ACM Press/Addison-Wesley Publishing Co., New York (2000)
18. Kovnatsky, A., Raviv, D., Bronstein, M.M., Bronstein, A.M., Kimmel, R.: Geometric and photometric data fusion in non-rigid shape analysis. *Numer. Math. Theory Methods Appl.* **6**, 199–222 (2013)
19. Krivokuca, M., Abdulla, W.H., Wunsche, B.C.: Sparse approximations of 3d mesh geometry using frames as overcomplete dictionaries. In: 2013 IEEE International Conference on Computer Vision Workshops (2013)
20. Lévy, B.: Laplace-beltrami eigenfunctions towards an algorithm that "understands" geometry. In: 2006 International Conference on Shape Modeling and Applications (SMI) (2006)
21. Maglo, A., Lavoué, G., Dupont, F., Hudelot, C.: 3d mesh compression: survey, comparisons, and emerging trends. *ACM Comput. Surv. (CSUR)* **47**, 44 (2015)
22. Mahadevan, S.: Adaptive mesh compression in 3d computer graphics using multiscale manifold learning. In: Proceedings of the 24th International Conference on Machine Learning. ACM (2007)
23. Mallat, S.G., Zhang, Z.: Matching pursuits with time-frequency dictionaries. *IEEE Trans. Signal Process.* **41**, 3397–3415 (1993)
24. Mohar, B., Alavi, Y., Chartrand, G., Oellermann, O.R.: The Laplacian spectrum of graphs. *Graph Theory Comb. Appl.* **2**, 12 (1991)
25. Neumann, T., Varanasi, K., Theobalt, C., Magnor, M., Wacker, M.: Compressed manifold modes for mesh processing. In: Computer Graphics Forum. Wiley Online Library (2014)
26. Pati, Y.C., Rezaifar, R., Krishnaprasad, P.S.: Orthogonal matching pursuit: recursive function approximation with applications to wavelet decomposition. In: 1993 Conference Record of The Twenty-Seventh Asilomar Conference on Signals, Systems and Computers, 1993. IEEE (1993)
27. Peng, J., Kim, C.S., Kuo, C.C.J.: Technologies for 3d mesh compression: a survey. *J. Vis. Commun. Image Represent.* **16**, 688–733 (2005)
28. Peter, P., Hoffmann, S., Nedwed, F., Hoeltgen, L., Weickert, J.: From optimised inpainting with linear pdes towards competitive image compression codecs. In: Pacific-Rim Symposium on Image and Video Technology. Springer, Berlin (2015)
29. Peyré, G., Mallat, S.: Surface compression with geometric bandelets. *ACM Trans. Graph. (TOG)* **74**, 601–608 (2005)
30. Pinkall, U., Polthier, K.: Computing discrete minimal surfaces and their conjugates. *Exp. Math.* **2**, 15–36 (1993)
31. Sorkine, O., Cohen-Or, D., Irony, D., Toledo, S.: Geometry-aware bases for shape approximation. *IEEE Trans. Vis. Comput. Graph.* **11**, 171–180 (2005)
32. Tosić, I., Frossard, P., Vanderghenst, P.: Progressive coding of 3-d objects based on overcomplete decompositions. *IEEE Trans. Circuits Syst. Video Technol.* **16**, 1338–1349 (2006)
33. Tropp, J.A., Gilbert, A.C., Strauss, M.J.: Algorithms for simultaneous sparse approximation. Part I: Greedy pursuit. *Signal Process.* **86**, 572–588 (2006)
34. Valette, S., Chaine, R., Prost, R.: Progressive lossless mesh compression via incremental parametric refinement. *Comput. Graph. Forum* **28**, 1301–1310 (2009)
35. Valette, S., Prost, R.: Wavelet-based progressive compression scheme for triangle meshes: wavemesh. *IEEE Trans. Vis. Comput. Graph.* **10**, 123–129 (2004)
36. Vallet, B., Lévy, B.: Spectral geometry processing with manifold harmonics. *Comput. Graph. Forum* **27**, 251–260 (2008)
37. Wallace, G.K.: The JPEG still picture compression standard. *ACM Commun.* **38**, 45 (1991)
38. Weickert, J.: Anisotropic Diffusion in Image Processing, vol. 1. Teubner Stuttgart, Stuttgart (1998)
39. Zhong, M., Qin, H.: Sparse approximation of 3d shapes via spectral graph wavelets. *Vis. Comput.* **30**, 751–761 (2014)



Yoni Choukroun received his BSc in Computer Engineering and his MSc at the Department of Computer Science, both from the Technion-Israel Institute of Technology. His research interests include computer vision and machine learning.



Gautam Pai is a PhD candidate at the Department Of Computer Science, Technion-Israel Institute of Technology. His research interests lie in the intersection of numerical geometry, signal processing, machine learning and computer vision.



Ron Kimmel is a professor of Computer Science at the Technion where he holds the Montreal Chair in Sciences and a part-time senior principal researcher in Intel RealSense. He held a postdoctoral position at UC Berkeley and a visiting professorship at Stanford University. He has worked in various areas of image and shape analysis in computer vision, image processing, and computer graphics. Kimmel's interest in recent years has been non-rigid shape processing and analysis, medical imaging and computational biometry, numerical optimization of problems with a geometric flavor, and applications of metric geometry and differential geometry. Kimmel is an IEEE Fellow for his contributions to image processing and non-rigid shape analysis. He is an author of two books, an editor of one, and an author of numerous articles. He is the founder of the Geometric Image Processing Lab. and a founder and advisor of several successful image processing and analysis companies.