Response to reviewer comments

Paper Number: 07C27
Author(s): Hao Zhang Oliver van Kaick Ramsay Dyer
Paper Title: Spectral mesh processing

Dear Editor and reviewers:

We greatly appreciate the thorough and thoughtful comments provided on our submitted article. It has taken us a rather long time to complete the final revision. We made sure that each one of the reviewer comments has been addressed carefully and the paper is revised accordingly.

As well, we have incorporated more images for better illustration of the concepts and added quite a few new references pertaining to the latest works on spectral geometry processing in the year 2008.

Attached below are detailed responses to all the reviewer's comments. The latter are shown in blue and our responses in red.

Please let us know if you still have any questions or concerns about the manuscript. We will be happy to address them, now in a timely manner.

Sincerely,

The authors of paper 07C27.

============
Review #1
============

Recommendation

Accept

Information for the Authors

Overall I find that this is a very nice survey, and a timely one, because there is no contemporary work that summarizes all the knowledge regarding spectral analysis of meshes and its applications in graphics. I have some minor comments regarding the manuscript and I hope the authors will consider addressing them, even though I recommended "accept" without revisions. Here are my suggestions in chronological order.

It would be better to formulate the abstract and the intro in present tense (i.e. "theoretical background is provided" instead of "will be provided").

>> This is done.

Sections 1 and 2 talk about the Laplacian without defining it. Surely, there are various definitions, as described later in the paper, but still it would be beneficial to give the reader some flavor or hint before plunging into the historical survey of section 2. Otherwise a novice reader will get lost quickly -- you talk so long about the spectrum of linear operators, but define one for the first time on page 7. Consider moving equation (1) and some of its related formalism to the beginning of the exposition.

>> This is done. The definition of the graph Laplacian is inserted into the start of Section 2. Also, a general description of the Laplace operator is provided at the same time. In addition, we have inserted two new figures to provide clear illustration of the spectral approach and spectral embedding early in the paper.

Section 2 is immediately followed by Section 4 (numbering needs correction).

>> We are not sure about this comment since there is indeed a Section 3: Overview of the Spectral Approach.

page 5, top: the text gives a (false) impression that any n x n matrix M has n real eigenvalues. Add that only certain types of matrices can be diagonalized, like the normal matrices over the field of complex numbers or symmetric matrices over the real numbers.

>> We could not locate the specified text, but suspect that this might be referring to the second paragraph of Section 5. This issue is now rectified.

page 8, very end: you say "The new operator does exhibit some nice properties. But a common drawback...". Please elaborate what "nice properties" mean in your context or remove this sentence (it's too vague). Same claim is repeated on page 9, top

(regarding T": When used in relevant applications, T" exhibits nice properties" - what applications and what properties?)

>> These issues are both addressed.

Somewhere in Section 6.2 you should mention the recent works by Wardetzky et al, namely:
"Discrete laplace operators: no free lunch". SGP 2007
"Discrete Quadratic Curvature Energies", CAGD 24, 2007.
These works have interesting and relevant discussion about the definition of the discrete Laplacian, the mass matrix and the related inner product on surfaces, symmetry, etc.

>> We added the two references and commented on them in Section 6.2.4 and Section 6.5.

Section 6.3.3: Q doesn't have a constant eigenvector because Q is constructed using the (correct) inner product on the discrete surface. This doesn't necessarily mean we can't use it for smoothing! Just that the flow needs to be constructed taking the right inner product into account.

>> Q is related to the Tutte Laplacian T. The latter operator requires the degree weighted inner product to render it self-adjoint and its eigenvectors orthogonal, however, Q is already symmetric. The zero row sum property expresses the fact that the constant vectors should lie in the kernel of the operator. If this is not the case, no choice of inner product can remedy the situation. In terms of the inadequacy of Q for smoothing, we are quite specific about the smoothing operation --- it is the low-pass filtering approach as done in [Taubin 1995].

Page 11: when you talk about the cotangent Laplacian, mention that it approximates the mean curvature normal (Mark Meyer's thesis can be cited, for instance).

>> We have elected to refrain from mentioning this interesting property of the Laplacian because it does not fit easily into the exposition and its relevance to spectral processing has not previously been expounded.

Page 18-19, Section 9.3: It is interesting to note that when a partial eigenbasis is used for projecting the geometry, we loose rotation-invariance of the representation in some sense: it is not the same thing to first rotate the mesh (hence changing its global coordinates) and then projecting it on a subset of the eigenbasis, or vice versa (first project and then rotate). At least it seems to me that projection and rotation operators don't commute.

>> The eigenprojection operation does commute with rotations; this can be seen as a consequence of the associativity of matrix multiplication. If P is the n x 3 matrix of vertex positions and R is a 3 x 3 rotation matrix, then $P = PR^T$ is the rotated coordinates. Let E be the n x n matrix of orthogonal eigenvectors and let E_k be the n x k matrix consisting of the first k columns of E. Then the eigenprojection of the vertices P using the first k eigenvectors is given by the k x 3 matrix $C = E_k^T P$. The rotation of

these projection coefficients yields the coefficients of the eigenprojection of the rotated geometry:

$$C \quad = CR^{T} = E_k{}^{T} P R^{T} = E_k{}^{T} P \quad .$$

This observation has been inserted into Section 8.3.

============
  Review #3
============

Recommendation

  Accept

Information for the Authors

Now that spectral mesh processing is doable (thanks to both hardware and software), this is the right time to publish this type of STAR. This is an important work, and a good reference (I especially like section 5, that gathers all useful theorems). Therefore, I recommend acceptance.

However, I've got some regret:

First, my philosophical two cents:
The way things are presented in this paper is completely discrete. The continuous setting is evocated at several places, but the relations between continuous and discrete are not explicited. I think that the explanations derived by Ramsay Dyer et.al in their tech report [DZM07] could be integrated in-extenso (in a section that would be called "relations between the continuous and the discrete setting"). This is the first time I've seen that clearly explained, so since R. Dyer is a co-author of this paper, I'm surprised not to see it there! Since most of the geometry-processing community does not "think in the continuous setting", I could do without it, but from my point of view, this would add some value in this paper.

>> The Laplacian operators are now treated in a separate section (Section 6) and the subsection on the geometric Laplacians (Section 6.5) has been rewritten to highlight the development of these operators via considerations of the properties of the differential operator, according to [DZM07].

Another place where I'd invoke the continuous setting is Section 4.1 when you say "it is commonly believed", I'd say that "it's true": If you compute "circular harmonics", i.e. the eigenfunctions of the second-order derivative, you exactly retrive the sines and cosines of the Fourier transform. Of course, this also works with a discretized circle, that has a circulant graph Laplacian, and that yields a discrete Fourier transform (as in Taubin's argumentation), but for me this is a "particular case" of the more general continuous setting. This is more obvious in the case of the eigenfunctions of Laplace-Beltrami on the sphere, you exactly retrieve spherical harmonics.

>> We acknowledge these facts pointed out by the reviewer and we added some of these to the paper. However, the original statement in the paper is meant to emphasize a resemblance between eigenanalysis with respect to the Laplace-Beltrami operator and classical Fourier analysis. There is only a resemblance as a key difference between the two situations exists; this is mentioned in the second paragraph of Section 4.1.

Second, I think that more structure could be added in the exposition (again, I can do without it: since the goal of this paper is to be exhaustive, this may be anti-nomic with presenting a nice coherent structure). For instance, historically, these methods were developed in the machine-learning community. An alternative way of structuring all these approaches would be to start from the following reference (I think this would give some tracks for answering the open questions listed at the end of the paper):

J. Ham, D. D. Lee, S. Mika, and B. Scholkopf, A kernel view of the dimensionality reduction of manifolds, technical report TR-110, Max Planck Institute for Biological Cybernectics (July 2003)

This one unifies Isomap, Laplacian Eigenmaps and LLE, as a special case of kernel PCA and also:

Y Bengio, P. Vincent, J.-F. Paiement, Spectral Clustering and Kernel PCA are Learning Eigenfunctions, 2003

This one unifies Kernel PCA and Spectral clustering as a special case of learning eigenfunctions (note the continuous setting)

The taxonomy is then as follows:

1. Learning Eigenfunctions
   - 1.1 Kernel PCA
     - 1.1.1 Isomap
     - 1.1.2 Laplacian eigenmaps
     - 1.1.3 LLE
   - 1.2 Spectral clustering

and I beleive that all the geometric methods can be attached to this hierarchy.

The "Kernel Trick" (http://en.wikipedia.org/wiki/Kernel_trick) can be used to explain these methods (for instance in the section about Gram matrices, that come from there)

The efficiency of these methods in the domain of geometry processing may be explained by:

Geometric Methods for Feature Extraction and Dimensional Reduction: A Guided Tour, Christopher J.C. Burges, November 2004, Technical Report: MSR-TR-2004-55

There were also recently several workshops on manifold learning and geometry:

http://nips.cc/Conferences/2007/Program/event.php?ID=597 http://people.cs.uchicago.edu/~niyogi/conference/speakers.html

... and I suspect that Kohonen's SOM (Self Organizing Maps) are a special case of those (i.e., a specialized numerical procedure to compute the second and third eigenfunctions of some operator)

(Keith Van Rijsbergen's "manifold" approach to Information Retrieval probably also belongs to that category of methods)

There are also very interesting references in the "References" section of the following website:

http://www.math.yale.edu/~sl349/tutorials.html

>> The reviewer has made some excellent points and we sincerely appreciate these well-thought comments. After considering these carefully, we decided not to introduce significant structural changes to the current paper, as the reviewer himself believes that these are not entirely necessary. However, we have added two paragraphs in Section 2.3 to allude to these relevant developments in the machine learning community and added the suggested references on kernel PCA, LLE, isomap, and Laplacian eigenmaps. As well, in quite a few other places in the paper, references to works from the machine learning communities are added to strength connections to this important relevant field.

More specific remarks follow:

Section 6.6.1: Gram matrices
I'd start by giving the general definition of a Gram matrix (i.e. matrix of the dot products of a familly of vectors), then particularize it to matrix learning. I'd mention the "Kernel trick" to justify the use of these matrices

>> Done. This is now Section 7.2.1.

Section 2, second paragraph: I'd add a reference to Belkin's paper about the convergence of graph Laplacians to their continuous counterparts.

>> Thank you for pointing this out. However, the graph Laplacian defined in the current paper is quite specific; it is what is usually known as the Kirchhoff operator --- it is combinatorial. Belkin抯 result applies to a graph Laplacian that is constructed by applying a Gaussian to the pair-wise Euclidean distances between points in a cloud. This is an important result and we mention this in Section 7.2.3, when we discuss non-sparse Laplacians.

Beginning of 2nd column: differential Laplacian coordinates: what is the link with spectral problems?

>> We do not see a clear link. This sentence is meant to give a historical account of the use of mesh Laplacians; works on differential coordinates are quite prominent in the geometry processing literature.

Section 3: overview of the spectral approach: M may be also obtained by discretizing a continuous operator.

>> Added in Section 3.

"Lagrangian of a graph" -> this one is intriguing (I do not know about that), could you elaborate on that?

>> This is defined in the paper now. See Section 3 after the bullet 揵ased on the operator used .

4.1: "Harmonic" behavior of Laplacian eigenvectors:
- please define "harmonic" (i.e. minimizer of Dirichlet energy)
- this can be simply explained by the min-max principle (theorem 5.2 or its continuous version)

>> Done.

4.3: "appropriately designed linear operator" with a specific application in mind, how would you "design" the operator?

>> Changed to 揷hosen .

Section 5: this is very nice to have all these theorems in the same place. Giving also a geometric intuition of what's going on could make this section even stronger.

>> Agreed.

End of Section 6.3.1: using 1/2 (T + T^t) as suggested by Levy is not very good. Vallet and Levy's tech report (07) gives a better solution.

>> We have now mentioned here that the use of a metric is a prefered way to treat T as a self adjoint operator. We also mention the work of Vallet and Levy in this context.

End of Section 6.3.2: Vallet and Levy's published a version of their tech report (Eurographics 08) with this point of view that may be interesting for the authors (that may also replace the citation at the send of 6.4)

>> We have develped a brief discussion of the DEC in Section 6.5.1 where we reference this derivation of the Laplacian presented in [Vallet and Levy 2008]. All refereneces to this work now refer to the Eurographics publication rather than the technical report.

Section 6.5: it could be interesting to explain what the continuous Schroedinger operator is (together with its original physical meaning)

>> Done. Now in Section 7.1.

6.7.1: non-sparse Laplacians: Belkin's discrete Laplacian is also non-sparse and could be discussed there.

>> Precisely. This is done now, Section 7.2.3.

============
Review #4
============

Recommendation

Accept after minor revision

Information for the Authors

The paper sums up nicely the latest advances in spectral mesh processing.

I had a few comments which might be helpful to the authors:

- A few more images of relevant applications could be useful: a graph showing the fast decay of spectral coefficients from [KG00], the image from [Rus07] showing the clustering of shapes with different deformations, segmentation images from [KLT05], the nodal sets and contours of eigenfunctions from [L06], etc.

>> Although we tend to agree, the applications listed in the survey are too many to choose from. Each will benefit from having an illustrative image, but we certainly cannot do that. After much debate, we have decided to be consistent and not show any image for the purpose of displaying results from these applications. We wish for the survey to serve as a single, comprehensive collection of pointers and a point of departure for readers to explore the relevant literature.

Having said that, we in fact have added new images to illustrate the main concepts underlying the spectral approach, e.g., an overview of the spectral approaching, construction of spectral embeddings, eigenvector plots, nodal line plots, etc.

- In the context of shape matching, the fact that there exist non-similar shapes which are isospectral, should be mentioned.

>> We added a reference to isospectral graphs and gave a citation, in Section 10.1.

- In section 6, it should be noted that the geometric Laplacian can be also be derived as the ratio of the dual edge length to the primal edge length, with different Laplacians emerging from different definitions of duality. This is derived and explained in [Gli05].

>> We have mentioned Glickenstein's work and duality structures in Section 6.5.1 We did not detail the derivation of the cotan operator from this perspective, but instead pointed to the exposition in [VL08,DHLM05] and]WBH*07].

- In section 6, the "No Free Lunch" theorem for Laplacians should be mentioned. [WMK*07].

>> This reference is added and commented in Section 6.2.4.

- The MDS section is a bit out of place in the operators section.

>> It is now in a subsection with a new title --- Section 7.2.3.

- The generalized eigenvalues problem could also be mentioned.

>> A new subsection has been added. -- Section 9.1

- In section 6, it should be mentioned that the combinatorial Laplacians are sensitive to the mesh triangulation.

>> A reference to the Connectivity Shapes paper has been added to emphasize that geometric properties captured by combinatorial Laplacians are lucky accidents. -- second paragraph introducing Section 6.

- It was nice, if Table 1 stated which operator and eigenstructure each paper used.

>> Perhaps. Though subsequent coverage in the section made these clear.

There are some theoretical questions which are not addressed by the authors, but I think are of interest:

- The mesh Laplacian can also be defined for meshes with boundaries. Does spectral analysis work as well in that case?

>> Generally yes. In this survey, we do not specifically address issues which might arise due to the existence of boundaries. Most techniques surveyed will work for both closed and open meshes, as long as they are 2-manifolds embedded in 3D. We have made clear of this at the start of Section 3.

- Is spectral analysis sensitive to small topologic changes - for example will two meshes with different genus but which are almost identical otherwise (for example, like the Homer mesh from [Rus07]) have similar eigenvectors and eigenfunctions?

>> For operators based on pairwise geodesic distances, the eigenstructures are not expected to remain stable even after small topological changes. However, for mesh Laplacians, the answer to this question is not entirely clear. We list this as a problem for future work.

References that should be added:

[WMK*07] Max Wardetzky, Saurabh Mathur, Felix K鍰berer, Eitan Grinspun.:Discrete Laplace operators: No free lunch. Symposium on Geometry Processing, 2007, pp. 33-37.

[Gli05] GLICKENSTEIN D.: Geometric triangulations and discrete Laplacians on manifolds. http://arxiv.org/math.MG/0508188. (2005).

>> These references have been added and discussed in the paper -- Sections 6.2.4 and 6.5.1