Title:
Adaptive Eigensystem Truncation for Spectral Shape Signatures

Authors:
Reed M. Williams, reed.williams@engr.uconn.edu, University of Connecticut
Horea T. Ilieş, ilies@engr.uconn.edu, University of Connecticut

Keywords:
Shape signatures, shape analysis, segmentation, HKS, spectral signatures, eigenvalues, SPCL

DOI: 10.14733/cadconfP.2016.274-278

Abstract:

The ability to compare the shapes of objects is crucial to the practice of engineering design. Spectral shape signatures provide a high-quality similarity measure based on diffusion physics by means of the spectrum of an estimate of the Laplace-Beltrami operator for the surface of an object.

However, point cloud and mesh models often have very large intrinsic sizes and subsequently large Laplace-Beltrami estimate matrices. Recommendations from the current spectral shape signature literature are to use only a low fixed number of arithmetically greatest eigenvalues and their corresponding eigenvectors in the computation of a spectral shape signature. This recommendation "seems to work well", but it is not yet understood the degree to which this fixed number of eigenpairs approximates the full spectrum for the purposes of shape similarity measures or even what fixed number to use. Using a fixed number of eigenpairs for all model sizes and samplings also introduces inconsistencies between different samplings of the same shape at different intrinsic sizes.

In this paper we examine the performance of fixed numbers of eigenpairs on approximating the spectrum of models of different sizes, propose an adaptive cutoff selection method which improves consistency between models for spectral signature use, and discuss the trade-off between running time for online applications and desired error or convergence properties.

Introduction and Background:

Understanding and comparing the shapes of parts and objects is fundamental to the design of functional structures. In the traditional of engineering practice, understanding shapes has been done intuitively by experts by means of their experience or by mathematical comparisons of simplified or representative shapes (e.g. combinations of primitives). In recent decades, methods have been developed for comparing shapes which do not rely on either disassembling shapes into representative primitives or engineering experience. Among the most useful for understanding three-dimensional shape is the class of techniques called spectral signatures.

A shape signature is a compact representation of shape which retains relevant information about the shape. A useful signature should retain enough information to discriminate completely between any two general shapes or classes of shapes while allowing straightforward computations of degree of similarity and remaining manageably sized. A particular class of shape signatures known as "spectral" shape signatures consists of signatures whose values are computed by reference to the spectrum of the Laplacian of a shape. The example spectral shape signature we select for our investigations here is the Heat Kernel Signature, although we remark that other spectral shape signatures (e.g., WKS) may be used instead.

A Fixed Number of Eigenpairs

Computing the values of a spectral shape signature on a shape requires computing the eigensystem of the Laplace-Beltrami estimate of that shape [5]. The Laplace-Beltrami operator for an *n*-point point cloud or *n*-vertex mesh model is an *n* x *n* matrix [2,4]. The complete eigensystem for such a model is *n* eigenvalues with *n* associated *n*-length eigenvectors. For typical CAD system or range scanner-generated models, *n* can easily be in the tens or hundreds of thousands or higher. Computing the complete eigensystem for a 200000x200000 matrix, even a sparse matrix, is an incredibly computationally intensive and time-consuming process.

In order to make a spectral signature for a typical model amenable to computation in a reasonable amount of time on a typical computer, the developers of spectral signatures traditionally have advised users to use a fixed number of eigenvalues and their associated eigenvectors for spectral signature calculation. For the Heat Kernel Signature, 300 eigenpairs is the recommendation [8]. For the Wave Kernel Signature, 300 is again the authors' preferred number [1]. The developers of the Global Point Signature used only 25 eigenvectors (though operating on decimated models of no more than 25000 vertices) [7]. The Shape Google implementation of HKS relies on only 100 eigenpairs [3]. In all of these cases the decision to compute only 25, 100, or 300 eigenpairs is justified only by experimental report that "it seemed to work well" for some test set.

In our corresponding journal paper, we characterize the level of approximation introduced to a typical spectral signature by fixed-number methods, discuss the limitations of these methods, and elaborate on our new tunable model-adaptive method of selecting the number of eigenpairs to use for each model in a database which helps mitigate those limitations. We also discuss tuning our method to adjust the balance between the computational speed and the precision of the computed signature.

Problem and Solution:

Spectral signatures are primarily used for shape similarity, comparing a query shape with the shapes of models in a database. Very rarely do all models in a database have a very similar intrinsic size. Often, even test databases, let alone real-world examples of the kinds of databases in daily use at engineering firms and manufacturing companies, contain models with an order of magnitude different numbers of vertices or points. For example, the CERTH/ITI Kinect scan database includes a model with 3657 points and another with 55808 points [11].

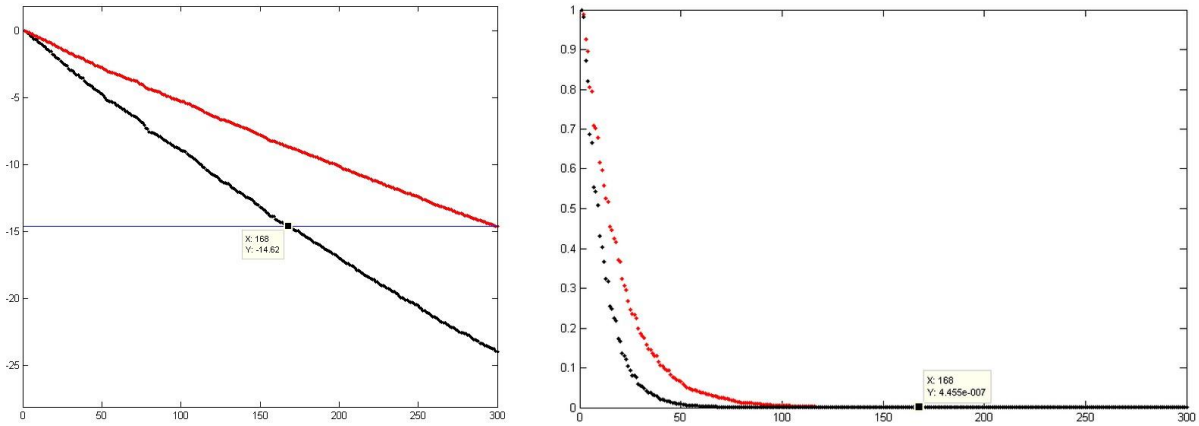

Fig: 1: The red points in each plot correspond to the larger 39k point robot model and the black points correspond to the smaller 26k point robot model. a) The first 300 eigenvalues of each model. The blue line marks the value of the 300th eigenvalue of the larger model and the labeled point is the first point below the line in the smaller model's eigenvalues. b) e^{λ} for the first 300 eigenvalues for each model. The label marks the final value in the larger model's plot.

Figure 1 shows the difference in the amount of information captured by the first 300 eigenvalues for two differently-sampled models of the same object. The eigenvalues are exponents in spectral signatures, so the lower values captured by the same fixed number of eigenvalues for the smaller model means capturing either information not present in the larger model's fixed number of

eigenpairs or, put differently, excessive information (and therefore excessive computational effort) if the larger model's amount of information is sufficient.

Limitations of Fixed-Number Methods

The primary limitations of these fixed-number methods are reduced precision for larger models, excessive computational effort for smaller models, and the introduction of a lack of consistency between measures which are supposed to be comparable.

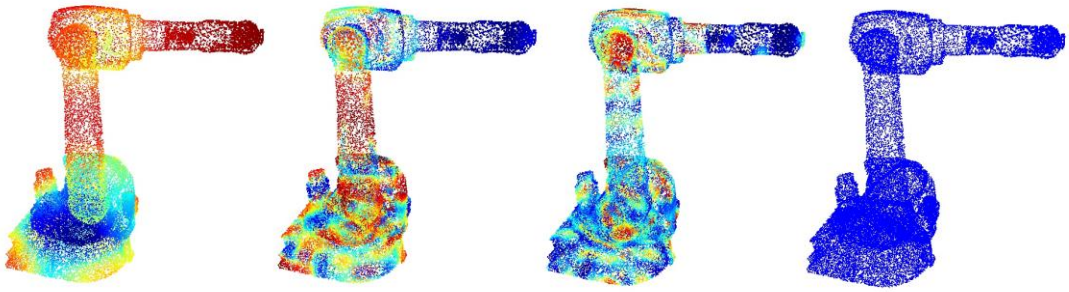

Fig. 2: a) The HKS vector at t = 0.03 (a model-based recommended large-scale t value) for the 26k robot model sampling, scaled to a unit bounding box, computed with 200 eigenpairs. The difference between this HKS vector as computed with different numbers of eigenpairs are shown in b)-d): b) shows the difference between HKS_{100} and HKS_{50}, c) the difference between HKS_{150} and HKS_{100}, and d) the difference between HKS_{200} and HKS_{150}. In these plots, red is higher differences and blue is lower. Note that the final plot, d), shows no difference at all between the computed HKS vectors.

In Figure 2, we offer an example of excessive computational effort which may be avoided by our tunable method. The figure shows the convergence of the t=0.03 HKS vector with respect to the number of eigenvalues and eigenvectors used in its computation is shown for a 26k point robot point cloud model. The HKS vector converges between 100 and 150 eigenpairs, as shown by the zero difference between the 150 and 200 eigenpair HKS vectors. This result implies that, for this model, sampling, and t-value, computing any more than 150 eigenpairs is wasted computational effort. For online processing, that extra effort and time can be the relevant factor in the timely detection of a feature or identification of an object.

Tunable Model-Adaptive Selection Method

Instead of computing a fixed number of eigenvalues, our method instead computes a small user-set number of eigenvalues and then predicts approximate successive eigenvalues $\hat{\lambda}$ by regression. These estimated eigenvalues provide a guide to the number of eigenpairs to compute for the spectral signature of the model in question. The estimated eigenvalues are examined for a point at which the contribution of the eigenvalue to the spectral signature is reduced below ξ (that is, $e^{t\lambda_n-1} - e^{t\lambda_n} < \xi$). This allows analysts to compute different numbers of eigenvalues for different models while capturing more similar portions of information encoded in those eigensystems. This procedure can be performed for any spectral signature scale (*t*-value), ideally the smallest *t*-value of interest for a given application (or $t = 1$). Once the estimated eigenvalues are examined, a point a short distance past the estimated location of the cutoff n is selected (we choose n+10 for its effectiveness in our test sets) and the eigenvalues and eigenvectors for the model are computed to that point. The new eigenvalues are checked to ensure that ξ has been reached; if it has not, the new set of eigenvalues just found are fed back into the quadratic estimator and the process begun again, using the additional information in the larger computed eigenvalues list to better guide the estimator.

The method (see Figure 3) is tunable mainly by two user-set parameters. The main tuning parameter is ξ, the cutoff difference. The cutoff difference specifies the minimum difference between pairs of subsequent eigenvalues. This parameter is where the majority of tuning for this method is performed. For our example in Figure 2, the value of $e^{\lambda,99} - e^{\lambda,100}$ was 120E-15 and the value of $e^{\lambda,149} - e^{\lambda,150}$ was 47E-21. The contribution to the HKS vector dropped more than a factor of a million across those fifty eigenvalues. The eigenvalue seed parameter we have fixed at 50 for convenience may

be adjusted based on performance on an analyst's system. The user can also choose a minimum and maximum number of eigenvalues to compute for any given model and the *t*-value to use for the database, based on their specific spectral signature and application.

Algorithm 1 Eigenvalue Cutoff Subroutine

```
function QTUNER(Laplacian,ξ,t,n_min,n_max)
    λ = eigs(Laplacian,50)
    while found ≠ TRUE do
        Fit quadratic to λ
        Use quadratic values as eigenvalue estimates to guess n from ξ
        ϕ, λ = eigs(Laplacian,n)
        gap = abs(exp(t·λ(1:end-1))-exp(t·λ(2:end)))
        n = min(max(min(where gap < ξ),n_min),n_max)
        if n == length(gap)and not length(gap) == n_max  then
            set found = TRUE
        else
            if length(gap) == n_max then set found = TRUE
            end if
        end if
    end while
    return λ,ξ,n
end function
```

Fig. 3: Our tunable method algorithm diagram.

Setting ξ = 100E-18 for the 14k point sampling of the robot model yields convergence at n = 113. For the 39k point sampling of that model, the same ξ setting yields convergence at n = 126, thirteen eigenpairs further than that at which the smaller model achieved the same degree of convergence and *significantly* less than the 300 eigenpairs recommended in [8], but more than the 100 recommended in [3]. The eigenvalues for each of these samplings are shown in Figure 4 down to the cutoff values recommended by our algorithm with ξ = 100E-18. Note that the eigenvalues reach approximately the same value despite the different numbers of eigenpairs computed.

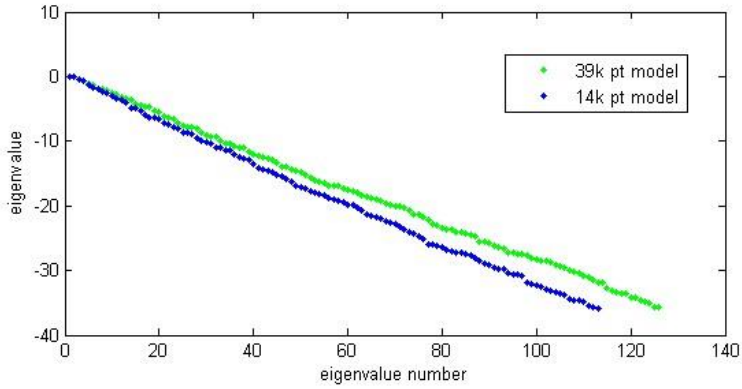

Fig. 4: The eigenvalues for the 14k and 39k point samplings of the robot model from the first eigenvalue to the n^{th} eigenvalue as recommended by our method.

Improved Consistency, Reduced Effort

This tunable adaptive cutoff method addresses the limitations of fixed-number methods discussed above. Computational effort is reduced while yielding the same effective amount of information. Much larger models (or spectral signatures at much smaller t-scale values) may require more eigenpairs than recommended by the fixed number methods to develop a similar level of convergence to smaller models in the same database. The enhanced consistency of the spectral signature result between models of different shapes and sizes allows greater user confidence in match candidates and segmentations based on spectral signature outputs.

Comparing the Top5 Hit Rate for HKS of the CERTH/ITI database using fixed 300 eigenpairs versus using the tunable adaptive cutoff method described above shows an improvement of 10% greater portion of same-category matches in the top five matches for each model. This demonstrates the importance of the enhanced consistency provided by the tunable method over a fixed number method in real application for matching and categorization.

Speed vs. Accuracy

We additionally note that while the example above demonstrated tuning to the degree that no further convergence of the signature vector was possible with additional eigenpair computation, the nature of the tunable method allows for intentional and well-understood under-convergence. That is, if additional speed is required for some online application or computational effort must remain limited (e.g., by hardware or power requirements), a cutoff ξ may be chosen to intentionally get only "enough" eigenpairs to allow, say, the degree of differentiation between shapes that your application requires. The tunable method allows this sort of designed "just enough" eigenpairs to be consistently specified across models of different sizes and over a range of scales.

Conclusion:

Spectral shape signatures are a popular class of similarity measures and have seen a great deal of use and many extensions in the literature in recent years. The number of Laplace-Beltrami eigenpairs used in computing these signatures was an important, but not yet well-understood, parameter. We have analyzed the limitations of fixed-number-of-eigenvalue methods for truncating model eigensystems for use in computing spectral shape signatures and developed a user-tunable method for adaptively determining required numbers of eigenpairs for different models of different shapes and scales.

Our tunable adaptive method improves consistency of signatures between models, enabling greater confidence in matching results and segmentations across large databases of different models from different scanning systems. This method can also greatly reduce computational overhead, especially valuable for systems where total computational power may be lower or resources may be in high demand, such as in autonomous systems (e.g. drones) or in real-time applications.

References:

[1] Aubry, M.; Schlickewei, U.; Cremers, D.: The wave kernel signature: A quantum mechanical approach to shape analysis, In Computer Vision Workshops (ICCV Workshops), 2011 IEEE International Conference on, 1626–1633. IEEE, 2011. http://dx.doi.org/10.1109/iccvw.2011.6130444

[2] Belkin, M.; Sun, J.; Wang, Y.: Discrete Laplace operator on meshed surfaces, In Proceedings of the twenty-fourth annual symposium on Computational geometry, 2008, 278–287. ACM. http://dx.doi.org/10.1145/1377676.1377725

[3] Bronstein, A.M.; Bronstein, M.M.; Guibas, L.J.; Ovsjanikov, M.: Shape google: Geometric words and expressions for invariant shape retrieval, ACM Transactions on Graphics (TOG), 30(1), 2011, 1. http://dx.doi.org/10.1145/1899404.1899405

[4] Evans, L.C.: Partial differential equations. The American Mathematical Society, 1998.

[5] Hsu, E.P.: Stochastic analysis on manifolds, 38. American Mathematical Society, 2002. http://dx.doi.org/10.1090/gsm/038

[6] Johnson, A.E.; Hebert, M.: Using spin images for efficient object recognition in cluttered 3d scenes, IEEE Transactions on Pattern Analysis and Machine Intelligence, 21(5), 1999, 433–449.

[7] Rustamov, R.M.: Laplace-Beltrami eigenfunctions for deformation invariant shape representation. In Proceedings of the fifth Eurographics symposium on Geometry processing, 2007, 225–233. Eurographics Association.

[8] Sun, J.; Ovsjanikov, M.; Guibas, L.: A concise and provably informative multi-scale signature based on heat diffusion, Computer Graphics Forum, 28, 2009, 1383–1392. Wiley Online Library, http://dx.doi.org/10.1111/j.1467-8659.2009.01515

[9] Williams, R.M.; Ilies, H.T.: Towards multi-scale heat kernel signatures for point cloud models of engineering artifacts, In Workshop on Algebraic Topology and Machine Learning at Neural Information Processing Symposium, 2012.

[10] The CERTH/ITI dataset of Kinect-based 3D scans. http://vcl.iti.gr/3d-scans/