Adaptive Eigensystem Truncation for Spectral Shape Signatures

Reed M. Williams[0000-0002-8192-6053][1] and Horea T. Ilieş[2]

[1]University of Connecticut, reed.williams@engr.uconn.edu
[2]University of Connecticut, ilies@engr.uconn.edu

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 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 and may cost unnecessary computational effort on resource-limited systems (e.g., drones or robots).

In this paper we briefly 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, demonstrate the method on Heat and Wave Kernel signatures (HKS and WKS) for point clouds, and briefly discuss the trade-off between running time and desired error or convergence properties.

Keywords: Shape signatures, shape analysis, segmentation, HKS, WKS, spectral signatures, eigenvalues, SPCL, point clouds.

1 INTRODUCTION

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.

1.1 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. Many of these signatures derive meaning by analogy to physical processes which are governed by the intrinsic geometry of the space in which they act. The example spectral shape signature we select for our investigations is the Heat Kernel Signature, although we remark that other spectral shape signatures may be used instead.

1.2 The Heat Kernel Signature

The Heat Kernel Signature (HKS) is a spectral shape signature based on the physical process of heat diffusion. It has a number of desirable properties: It is invariant up to model isometry, intrinsically multi-scale, and stable under perturbations on the scale of typical range camera noise. In order to get a physical sense for the meaning of the HKS of a shape, consider a point source of heat applied to a point on a surface. As time passes, the heat will diffuse on the surface away from that point. The heat kernel signature's value k_t at that point is the sum total of all of the heat which has diffused away by time t. Since the Laplacian describes the flux of a vector field on a surface, it is intuitive that the Laplace operator will be of some use in this computation. Indeed, the heat equation on a manifold is defined as

$$\frac{\partial u}{\partial t} - \alpha \nabla_M u = 0 \tag{1}$$

where α is a positive constant and u is the thermal energy as a function of time and location on the surface [6].

The heat kernel is a fundamental solution to the general heat equation [7]. The heat kernel in terms of the eigensystem of the Laplace-Beltrami operator may be written as

$$k_t(x,y) = \sum e^{-\lambda_i t} \phi_i(x) \phi_i(y) \tag{2}$$

The quantity $k_t(x,y)$ may be considered equivalent to a measurement, for time t, of the amount of heat transferred from point x to point y, for some initial distribution of heat energy on the surface $u_0(x)$. Using this quantity as a measure for similarity would require mappings between each of the neighborhoods, mappings which would be difficult or time-consuming to define between models.

The heat kernel signature (HKS) of a shape is a more compact description of a shape than the heat kernel itself; it is a restriction of the heat kernel to $k_t(x,x)$, the diagonal of the heat kernel. This restriction captures the “amount of heat” which has diffused away from point x by “time” t. This restricted version of the heat kernel is sufficient to describe the local area of point x for the purposes of similarity. Restricting the heat kernel to the “time” domain over the model reduces the computational complexity of the signature and obviates the need to develop these local mappings for similarity. This form of the heat kernel on M has the eigendecomposition

$$k_t(x,x)_M = \sum e^{-\lambda_i t} \phi_i(x) \phi_i(x) \tag{3}$$

where λ and ϕ are the eigenvalues and eigenvectors of the Laplace-Beltrami operator of M. Thus, the first step in calculating the HKS for a given shape must be the estimation of the Laplace-Beltrami operator of the surface of the shape.

1.3 Laplace-Beltrami Estimate

The Laplace operator or Laplacian is a second-order differential operator Δf which describes the variation of a differentiable function f within a space. It is defined as the divergence of the gradient of the function

$$\Delta f = \nabla \cdot \nabla f \tag{4}$$

which is equivalent to the sum of the unmixed second-order partial derivatives. Intuitively, this describes the flux of the gradient field of a function in that space. The equivalent form on a Riemannian (i.e., real, smooth, equipped with an inner product) manifold is called the Laplace-Beltrami operator.

$$\Delta_M f = \mathrm{tr}(\mathrm{H}(f)) \tag{5}$$

The Hessian H(*f*) of the function is a square matrix of second-order partial derivatives that describes the local curvature of the function f over the manifold. Taking the trace of the Hessian keeps only the unmixed second derivatives, as in the definition of the standard Laplace operator.

This description of local curvature makes the Laplace-Beltrami operator of a surface a powerful tool for shape analysis. Discretizations of the Laplace-Beltrami operator for various discrete representations of a surface have been the subject of intense academic interest [12]. For triangular surface meshes, the current state-of-the-art is the Mesh Laplace Operator [2]. For point cloud models, the equivalent estimate is the Symmetric Point Cloud Laplacian [13-14].

1.4 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 (see Equation 3). The Laplace-Beltrami operator for an n point or n vertex model is an $n \times n$ matrix. 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. It has been suggested and broadly accepted that the “rapid convergence” of the eigenvalues should allow for a signature to be computed “using a moderate number of eigenvalues … determined by feasible numerical computations" [9].

Thus, in order to make a spectral signature for a typical model amenable to computation, the developers of spectral signatures have traditionally 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 [11]. 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) [10]. The Shape Google implementation of HKS relies on only 100 eigenpairs [4]. 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. How many eigenpairs should be computed to allow the “rapid convergence” of the spectrum to converge appropriately for any given model is still an open question. Other works have suggested that some other subset of the spectrum provides better discriminatory power for particular cases, but ultimately conclude that the “first k eigenvalues” methods perform at least as well as those alternative subsets in general (see Table 2 in [8]). Additionally, alternative subsets take significant offline processing to develop for a given database [8].

1.5 In This Paper

In this work, 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 present analysis and discussion of tuning our method to adjust the balance between the computational speed and the precision of the computed signature. Although we use point cloud models and the SPCL estimate [14] of the Laplace-Beltrami operator throughout the experiments, the approach presented in this paper can be applied to compute spectral signatures for other types of models as long as a convergent estimate of the Laplace-Beltrami operator exists.

2 UNDERSTANDING THE IMPACT OF EIGENPAIR CUTOFF

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 contain models with an order of magnitude difference in the number of vertices or points, let alone real-world examples of the kinds of databases in daily use at engineering firms and manufacturing companies. For example, the CERTH/ITI Kinect scan database includes a model with 3657 points and another with 55808 points [15].

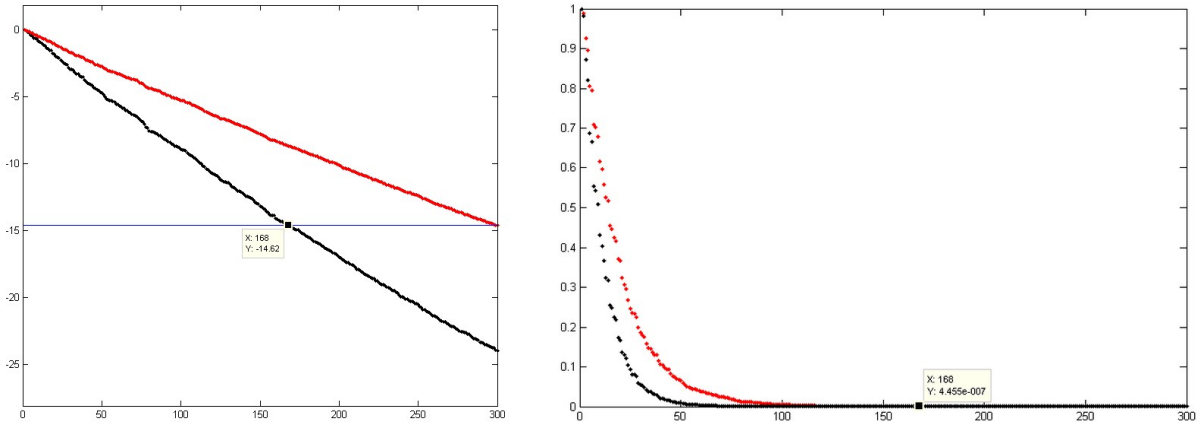

Fig. 1. The red points in each plot correspond to the larger 39k point robot model and the black points correspond to the smaller 26k point robot model. a) The first 300 eigenvalues of each model. The blue line marks the value of the 300th eigenvalue of the larger model and the labeled point (168) is 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 using N eigenvalues for the smaller model means capturing information not present in the larger model's N eigenpairs. Put differently, this means capturing excessive information (and therefore using excessive computational effort) if the larger model's amount of information is sufficient to the application. Two different samplings of the same model are used here rather than two different models of different sizes purely for clarity. The analogy of sufficient information holds even across models of different shapes.

2.1 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.

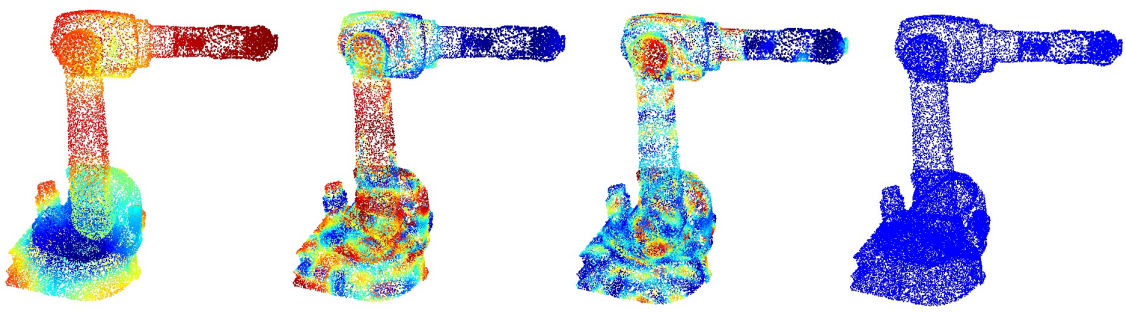

a) b) c) d)

Fig. 2. a) The HKS vector $k_{0.03}$ 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: The vectors have converged somewhere between 100 and 150 eigenpairs of information for this model at this t-scale.

In Figure 2, we offer an example of the kind of excessive computational effort which the method we introduce avoids without loss of precision. The figure shows the convergence of a particular 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 and sampling, at this t-value, computing any more than 150 eigenpairs is wasted computational effort. For online processing, that extra effort and the time associated with it can be the relevant factor in the timely detection of a feature or identification of an object.

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_0 from ξ
        φ, λ = eigs(Laplacian,n_0 + 10)
        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.

3 TUNABLE MODEL-ADAPTIVE SELECTION METHOD

Instead of computing a fixed number of eigenvalues, our method instead computes a quite small user-set number of eigenvalues and then predicts approximate successive eigenvalues $\tilde{\lambda}$ by regression. These estimated eigenvalues provide a guide to what 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 in question to the spectral signature is reduced below a parameter 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 the information encoded in those eigensystems.

This procedure can be performed for a given spectral signature scale (*t*-value), ideally the smallest *t*-value of interest for a given application, or with $t = 1$. 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 to avoid underestimating n) and the eigenvalues and eigenvectors for the model are computed up to that location. 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. This should nearly always result in reaching ξ in a maximum of two eigensystem computations after the first 50-length computation.

The method (see Figure 3) is tunable mainly by two user-set parameters. The main tuning parameter is ξ, the cutoff difference, which specifies the minimum difference between pairs of subsequent eigenvalues. This parameter controls most of the tuning performed by this method. 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. This parameter allows the user to choose how many eigenvalues to compute before fitting the quadratic and predicting the convergence of $e^{\tilde{\lambda}}$. 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.

We note that spectral signatures hypothesize as part of their formulation that Laplace-Beltrami spectra will not contain repeated eigenvalues [11]. As there is no general understanding of the operation of spectral signatures outside that restriction, we do not concern ourselves with the possibility of high-multiplicity eigenvalues appearing at a critical juncture and interfering with the cutoff computation.

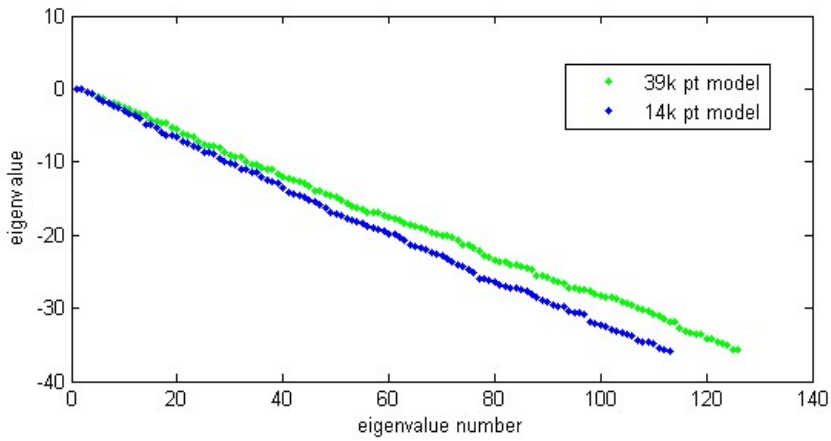

Fig. 4. The value of eigenvalues for the 14k and 39k point samplings of the robot model from the first to the cutoff number recommended by our method. Note the similar final values despite the different numbers of total eigenvalues computed to reach that point.

Setting ξ = 100E-18 for the 14k point sampling of the robot model yields convergence at $n = 113$. For the 39k point sampling of the same robot 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 [11], but more than the 100 recommended in [4]. 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 real value despite requiring different numbers of eigenvalues to be computed to reach that value.

3.1 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. This permits more efficient signature development without effective loss of precision. Much larger or more complex models 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 matching candidates and segmentations based on spectral signature outputs. Sufficient inconsistency between spectral fraction used to compute a shape signature between two models in a database may lead to misidentification or misclassification of shapes. This technique helps to avoid such inconsistency.

3.2 Tuning For Speed vs. Precision

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 the degree of differentiation between shapes that your application requires. The tunable method allows this sort of designed "just enough" quantity of eigenpairs to be consistently specified across models of different sizes and over a range of scales.

As well, eigensystem computation speed does not scale linearly, so even if computing a consistent number of eigenpairs takes two or even three calls to the eigenvalue solver, so long as the average final number of eigenpairs is lower than would be chosen by a fixed-number method the total time to compute the spectral signature will be less. We note that nearly every model in the example in Section 4 reached ξ in only two calls to the eigenvalue solver (including the initial seed call of only 50 values). For example, in the following demonstration, for the first model of the CERTH/ITI database, the tunable cutoff method eigensystem call takes less than half the time of a traditional 300-pairs eigensystem call.

4 FURTHER DEMONSTRATION

In order to demonstrate the use of our method in a general case for real scanner data, we present the following example on the CERTH/ITI Range Scan Dataset [15], a freely-available database of scanned point cloud models of a variety of small objects produced by a Microsoft Kinect sensor with a depth resolution of about 1 cm [5]. See Figure 5 for photographs of a sample of the objects from the database.

The objects in the database were scanned in eighteen different rotations of a turntable. The database provides an .XYZ file of the set of scans of each object rotated into a common coordinate system (so-called "registered" scans). Minimal cleaning has been performed to remove points outside the bounds of the turntable (i.e., background removal), but outlier points remain, and the points of the aligned scans are often positioned so that a reconstructed surface through the points would result in self-intersections and other surface degeneracies. These scans, which are similar to those produced by industrial and hobby range scan systems, would be challenging to mesh without human operator intervention. We use the point cloud model HKS analysis method of [14] to develop HKS vectors from the models and we *a priori* select minimum n = 80, maximum n = 500, and cutoff ξ = 100E-12. Figure 6 shows a handful of similar models from the database colored by those HKS vectors.

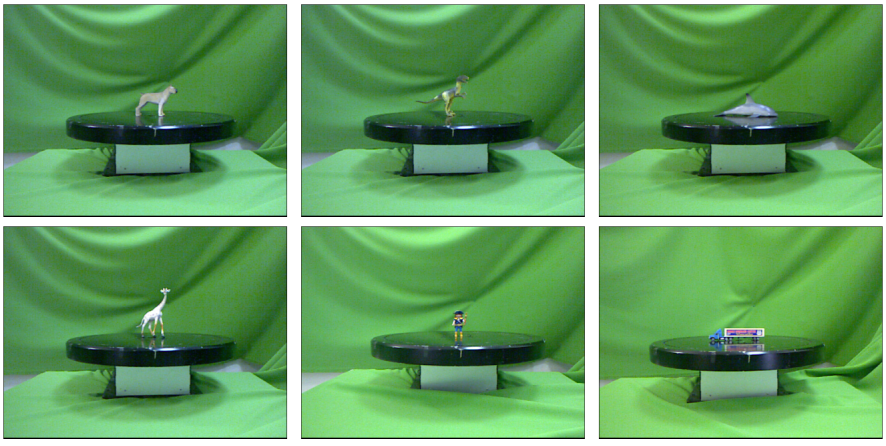

Fig. 5. A sample of the kinds of shapes present in the CERTH/ITI database.

The tunable model-adaptive selection method suggested a variety of eigenvalue truncation levels for the CERTH/ITI Kinect model database. The minimum size suggested was 80 eigenpairs for a model of 5046 points while the maximum suggestion was limited by an a priori maximum of 500 eigenpairs. Figure 7 shows the number of eigenpairs used in computing the heat kernel signature vectors for each model plotted against model size.

Additionally, 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 a marked 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.

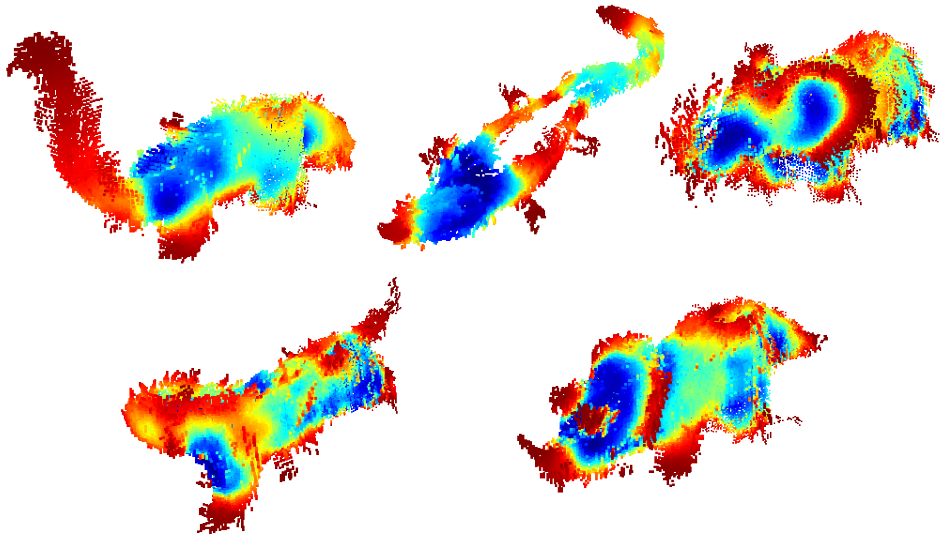

Fig. 6. A sample of the point cloud models from the CERTH/ITI database showing heat kernel signature vectors computed using the truncated eigenvalues and eigenvectors suggested by the tunable adaptive cutoff method.

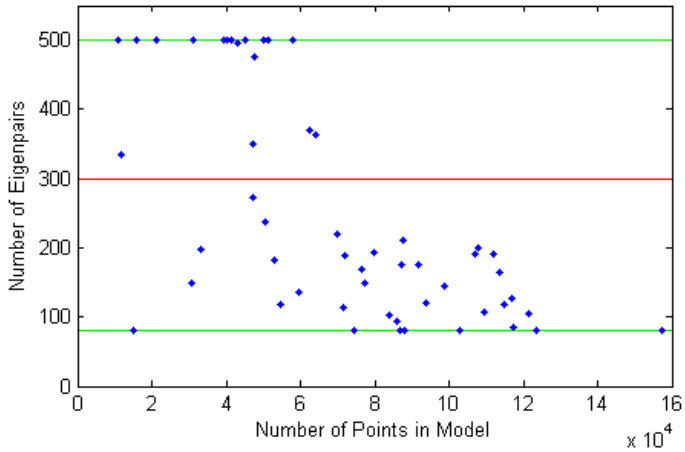

Fig. 7. Number of eigenpairs suggested (limited to a minimum of 80 and a maximum of 500 eigenpairs) by the tunable model-adaptive cutoff method for the models in the CERTH/ITI Kinect scan database plotted by number of points in the (automatically) trimmed model. The red line at 300 represents the canonical suggested number of eigenpairs for HKS for each the model. For any model above the line, 300 values underrepresents the Laplacian reducing consistency and for any model below the line, 300 values represents more computational effort or over-specificity. Note that the number of eigenpairs is not determined by model size.

4.1 Wave Kernel Signature

To demonstrate that our methods are applicable to spectral signatures other than the Heat Kernel Signature, which we use for our other examples and explanations in this manuscript, we demonstrate the usability of the proposed approach on the Wave Kernel Signature [1] as well. In Figure 8, we show a segmentation of a noisy robot model by the Wave Kernel Signature using a) 300 eigenpairs and b) an automatically-computed 222 eigenpairs chosen by our tunable adaptive selection method (using the same ξ and min and max ns as in the CERTH/ITI example). Note the similarity of the segmentations. Because the spectral signature has converged by 222 eigenvalues, the additional eigenvalues available to the n=300 signature do not change the segmentation significantly. Additionally, the adaptively-tuned model required ~7% less run time.

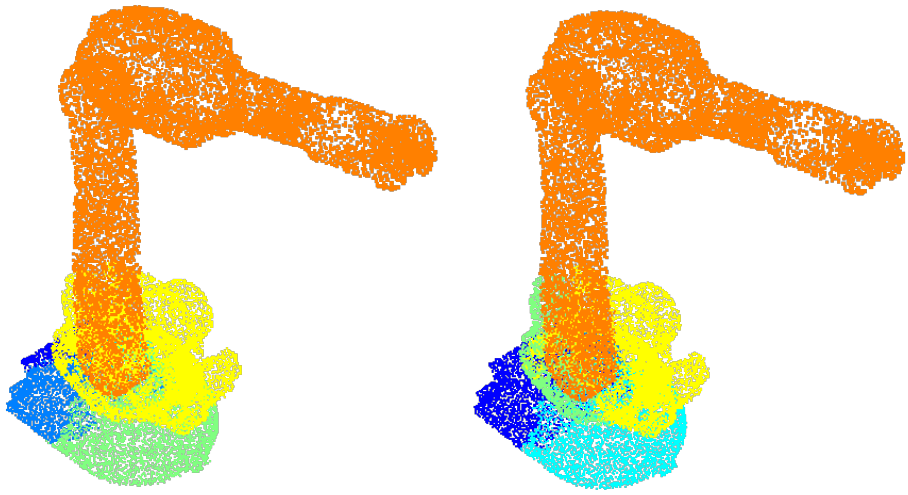

a) b)

Fig. 10. The ~26000 point robot model segmented based on the WKS at e = -1.4756 using a) a fixed 300 eigenpairs as recommended in the original paper and b) using an automatically-computed 222 eigenpairs chosen by our tunable model-adaptive method. Note that colors are randomly assigned to segments and similarity of color should not be taken to mean similarity of segment between or within model.

5 SUMMARY AND CONCLUSIONS

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 discussed 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 between models of different samplings, 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, enabling online use of techniques on 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. Because this method is a modification to the general method of spectral shape signatures, the benefits obtained thereby can be combined with the advantages of any present or forthcoming published enhancements to spectral signature technology.

Once spectral signatures have been generated for a particular database, integrating additional scans may be made even more efficient by developing a function mapper to guess the cutoff number of eigenvalues for the new scan from information about the models which are already in the database. Such a function mapper could be a neural network or regression that develops an estimate of eigenvalue cutoff number n from implicit and explicit sizes of models already analyzed. This could enable closer-to-realtime scan analysis by reducing computational overhead for mesh or point cloud models. Another recent paper proposed a method for understanding which functions (e.g., signature values) are most discriminative for a shape [3]. Applying their methods to co-segment a database with signature values computed with enhanced inter-model consistency by our tunable modification may produce a high quality database segmentation quickly and consistently.

ACKNOWLEDGEMENTS

This work was supported in part by the National Science Foundation grants CMMI-1200089, CMMI-0927105, and CNS-0927105. Reed Williams was also partially supported through the General Electric Fellowship for Innovation.

Reed M. Williams, http://orcid.org/0000-0002-8192-6053

REFERENCES

[1] Aubry, M.; Schlickewei, U.; Cremers, D.: The Wave Kernel Signature: A Quantum Mechanical Approach To Shape Analysis, 2011 IEEE International Conference on Computer Vision Workshops, 2011, 1626–1633. http://dx.doi.org/10.1109/iccvw.2011.6130444

[2] Belkin, M.; Sun, J.; Wang, Y.: Discrete Laplace operator on meshed surfaces, Proceedings of the 24th Annual Symposium on Computational Geometry, 2008, 278–287. http://dx.doi.org/10.1145/1377676.1377725

[3] Biasotti, S.; Spagnuolo, M.; Falcidieno, B.: Grouping real functions defined on 3D surfaces, Computers & Graphics, 37(6), 2013, 608-619. http://dx.doi.org/10.1016/j.cag.2013.05.007.

[4] 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, 1-20. http://dx.doi.org/10.1145/1899404.1899405

[5] A. Doumanoglou, S. Asteriadis, D. Alexiadis, D. Zarpalas, P. Daras, “A Dataset of Kinect-based 3D scans”, 11th IEEE IVMSP Workshop: 3D Image/Video Technologies and Applications, Yonsei University, Seoul, Korea, 2013.

[6] Evans, L.C.: Partial Differential Equations, American Mathematical Society, Providence, RI, 1998.

[7] Hsu, E.P.: Stochastic Analysis on Manifolds, American Mathematical Society, Graduate Studies in Mathematics series (Book 38), Providence, RI, 2002.

[8] Marini, S.; Patané, G.; Spagnuolo, M.; Falcidieno, B.: Spectral feature selection for shape characterization and classification, The Visual Computer, 27(11), 2011, 1005-1019.

[9] Reuter, M.; Wolter, F.; Peinecke, N.: Laplace-Beltrami spectra as "Shape-DNA" of surfaces and solids, Computer-Aided Design, 38, 2006, 342–366. http://dx.doi.org/10.1016/j.cad.2005.10.011

[10] Rustamov, R.M.: Laplace-Beltrami eigenfunctions for deformation invariant shape representation, Proceedings of the fifth Eurographics symposium on Geometry processing, Eurographics Association, 2007, 225–233. http://dx.doi.org/10.1109/34.765655

[11] 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. http://dx.doi.org/10.1111/j.1467-8659.2009.01515

[12] Wardetzky, M.; Mathur, S.; Kalberer, F.; Grinspun, E.: Discrete Laplace Operators: No Free Lunch, Symposium on Geometry Processing, 2007, 33-37.

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

[14] Williams, R.M.; Ilies, H.T.: Practical Shape Analysis and Segmentation Methods for Point Cloud Models. Technical report, preprint, 2016.

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