

Feature Selection for Enhanced Spectral Shape Comparison

S. Marini, G. Patané, M. Spagnuolo, B. Falcidieno

Consiglio Nazionale delle Ricerche,
Istituto di Matematica Applicata e Tecnologie Informatiche
Genova, Italy

Abstract

In the context of shape matching, this paper proposes a framework for selecting the Laplacian eigenvalues of 3D shapes that are more relevant for shape comparison and classification. Three approaches are compared to identify a specific set of eigenvalues such that they maximise the retrieval and/or the classification performance on the input benchmark data set: the first k eigenvalues, by varying k over the cardinality of the spectrum; the Hill Climbing technique; and the AdaBoost algorithm. In this way, we demonstrate that the information coded by the whole spectrum is unnecessary and we improve the shape matching results using only a set of selected eigenvalues. Finally, we test the efficacy of the selected eigenvalues by coupling shape classification and retrieval.

Categories and Subject Descriptors (according to ACM CCS): I.3.5 [Computer Graphics]: Computational Geometry and Object Modelling— Additional keywords: *Shape comparison and matching, Laplacian spectrum, feature selection, AdaBoost.*

1. Introduction

Shape classification and retrieval are crucial tools to organise and interact with databases of 3D models and to get a picture on the knowledge, or semantics, underlying the models. The performance of classification and retrieval strongly depends on the effectiveness of the shape descriptors, the comparison method, and the indexing techniques [BFF*07]. While several methods have been proposed to compare 3D shapes [BKS*05, TV08], only few methodologies address the issue of identifying descriptions that capture the shape features shared by a class of models [MSF07, LN08].

Inspired by the earlier work presented in [TV04, LN07], the goal of this paper is to deepen the analysis of the classification and retrieval performances of feature vectors defined by the spectrum of the Laplace-Beltrami operator. The result of the study is a new approach to automatically associate to classes of 3D objects the sub-set of the spectrum that is more relevant to characterise the inter-class similarity and discriminate among different classes. To identify a specific set of eigenvalues that maximises the retrieval and/or the classification performance on the input benchmark data set, we have investigated three approaches: the first k eigenvalues, the Hill Climbing technique, and the AdaBoost algorithm. The choice of these three algorithms for feature selection is

motivated by the need of a statistical approach to the definition of a *feature space*, where close points correspond to shapes with similar characteristics within the same class.

Even if the aforementioned algorithms are not new, our work is the first attempt at the identification of those sub-parts of the shape spectrum which discriminate among different classes of models. More precisely, our final aim is to identify clustered eigenvalues, which describe the associated class by means of those local/global features that are discriminative for that class. The AdaBoost algorithm [FS99], which is a statistical tool for feature extraction, has been recently used for 2D images [TV04] and for 3D models [LN07] to select relevant views of 3D objects with respect to the light field descriptor [CTSO03]. Other classifiers based on semi-supervised learning, dimensionality reduction, and probability have been successfully exploited for shape classification. For instance, in [HLR05] Support Vector Machine is used to cluster 3D models with respect to semantic information. In [HR06, OK06], shape classifiers are obtained as a linear combination of individual classifiers and using non-linear dimensionality reduction. In [SF06], relevant local shape descriptors are selected through a multivariate Gaussian distribution and collected to define a priority-driven search for shape retrieval.

In the context of 3D shape analysis and matching, the spectrum of the Laplace-Beltrami operator provides a descriptive and large feature vector, which characterises the input shape and has been widely studied in the recent years. We refer the reader to [ZvKD07] for a recent survey on spectral mesh processing. The Laplacian spectrum is well suited for shape matching tasks due to its isometry-invariance properties, its robustness to local noise and sampling, and its shape-intrinsic definition and multi-scale organisation. The first use of the Laplacian spectrum for shape matching was firstly proposed in [RWP06]. Here, two shapes are compared by measuring the distance between the vectors defined by the first $k := 50$ eigenvalues with smallest magnitude.

Instead of using the spectrum itself, in [JZ07] non-rigid objects are matched using spectral embeddings, which are derived from the eigenvectors of affinity matrices computed considering geodesic distances. In [Rus07], the first k Laplacian eigenvalues and eigenvectors are used to define an isometry-invariant shape representation. Then, these signatures are compared using a modification of the $D2$ -distribution [OFCD02], which is based on a set of histograms that capture the variation of distances among points within a set of spherical cells centred at the origin of a k -dimensional space. In [MCHB07], spectral embeddings are constructed as *Local Linear Embedding* on eigenspaces of affinity matrices and matched by using the *Expectation-Maximisation* framework. Finally, the Laplace-Beltrami operator is strictly related to the heat diffusion equation, which provides an embedding of a given scalar function in a hierarchy of smoothed approximations. Recently, the heat equation and the associated diffusion metric have been used to define multi-scale shape signatures [SOG09], compare 3D shapes [M09], and approximate the gradients of maps defined on triangulated surfaces and point sets [Wan09].

While there is an evidence of the close relationship among shape features and eigenvalues, the best way to use the spectrum for shape characterisation has not been identified yet [RBG*09]. We argue that statistical methods are the most appropriate to correlate sub-sets of the spectrum to classes of 3D shapes and to have a grasp on the *semantics* captured by the eigenvalues. In this context, our main contribution is the application and comparison of three feature selection approaches to the Laplacian eigenvalues (Section 3); namely, (i) the *first k eigenvalues*, by varying k on the cardinality of the computed spectrum; (ii) the Hill Climbing technique; and (iii) AdaBoost [FS95]. The obtained results confirm the hypothesis that the Laplacian spectrum contains unnecessary information for shape matching and classification. Indeed, the appropriate selection of a set of eigenvalues strongly improves the classification results and the retrieval efficacy (Section 4). Finally (Section 5), we provide closing remarks on results and outline future work.

2. Spectral descriptors of 3D shapes

The intuition behind spectral shape comparison is that the Laplacian spectrum is meaningful for the description of the input surface \mathcal{M} due to its intrinsic definition, invariance with respect to isometric transformations, and easy computation. More precisely, the *Laplacian spectrum* of \mathcal{M} is defined as the set of solutions (λ, f) of the following *eigenvalue problem*: find $f : \mathcal{M} \rightarrow \mathbb{R}$ such that $\Delta f = \lambda f$, $\lambda \in \mathbb{R}$, where Δ is the Laplace-Beltrami operator. In the discrete setting, let us consider a triangulated surface \mathcal{M} with vertices $\{\mathbf{p}_i\}_{i=1}^n$. Then [RWP06], the FEM discretisation is equivalent to the *generalised eigenvalue problem* $L_{\text{cot}}\mathbf{f} = \lambda B\mathbf{f}$, $\mathbf{f} := (f(\mathbf{p}_i))_{i=1}^n$, where the $n \times n$ matrices L_{cot} and B are the stiffness matrix with cotangent weights and the mass matrix, respectively. Alternatively, we can compute the eigensystem associated to the *standard eigenvalue problem* $(\lambda, \mathbf{f}) : L_{\text{cot}}\mathbf{f} = \lambda \mathbf{f}$. In the following, we assume that the eigenvalues λ_k and the corresponding eigenvectors \mathbf{f}_k , $k = 1, \dots, n$, are increasingly reordered; i.e., $0 = \lambda_1 \leq \lambda_2 \leq \dots \leq \lambda_n$.

Normalisation of the FEM Laplacian spectrum. For shape comparison, the main properties of the spectrum of the Laplacian matrix of a 3D surface are its isometry-invariance, which does not require shape alignment, and its independence of discretisations. Since uniformly rescaling \mathcal{M} by a constant factor s changes the FEM Laplacian eigenvalues by s^{-2} , we make them invariant to shape scales by normalising the spectrum of \mathcal{M} with its area. Furthermore, as k tends to infinity, λ_k becomes close to the value $4\pi k / \text{area}(\mathcal{M})$. Note that the eigenvalues of L_{cot} are not affected by a rescaling of \mathcal{M} and are bounded from above. Since the FEM spectrum is unbounded from above, for our experiments we normalise the computed eigenvalues by their maximum.

Assuming that the sampling density of the input surface is coherent with the shape details that must discriminate similar shapes, our experiments have shown that the normalised spectrum of the Laplacian matrix, discretised with both the cotangent and FEM weights, is not strongly affected by the noise and shape discretisation in terms of connectivity and sampling. For instance, in the examples of Fig. 1 smoothed, simplified, and isometrically deformed surfaces have an almost identical spectrum. Theoretical and experimental results on the sensibility of the eigenvalue computation with respect to noise, sampling, and deformation have been discussed in [DZMC07, RBG*09]. Even though the Laplacian spectrum characterises geometric and topological features of 3D shapes in a way that is not unique, previous work [RWP06, RBG*09] has also shown that the spectrum is able to distinguish dissimilar shapes.

3. Classification based on the shape spectrum

In the literature, few works tackle the problem of similarity by using the shape spectrum. Even though the approach pro-




Figure 1: Variation (y-axis) of the first 100 Laplacian eigenvalues (x-axis) on isometrically deformed shapes: (a) FEM and (b) cotangent weights; the FEM eigenvalues have been normalised with respect to the surface area and their maximum. (c-e) On the left and right side, deformed (50K vertices) and simplified (10K vertices) shapes used to verify the sensibility the eigenvalues’ computation.

posed in [RWP06] provides good results, it is not clear if the selection of a particular spectrum sequence is the best choice or there exist other sequences of eigenvalues that provide better results. Furthermore, the large amount of extracted information stresses the importance of identifying relevant information from the shape spectrum.

In this context, our work investigates the problem of selecting a bunch of eigenvalues that are capable of maximising the performances of classification and retrieval algorithms. To this end, we compare three approaches to select a specific set of eigenvalues such that the corresponding shape classification error on the input benchmark data set is minimised: the *first k eigenvalues* (Section 3.1), by varying k on the cardinality of the computed spectrum; the Hill climbing algorithm (Section 3.2); and the *AdaBoost algorithm* (Section 3.3). Using the aforementioned methods, the shape classification has been investigated by selecting those sub-sequences of the Laplacian spectrum that characterise the inter-class similarity and discriminate among different shape classes. The distance \tilde{d}_s between the query \mathcal{Q} and the class C is defined as

$$\tilde{d}_s(q, C) = \min_{m \in C} d_s(q, m), \quad (1)$$

where d_s is the distance between the models with respect to s . The following is a general definition for the query-class classification scheme

$$q \mapsto \bar{C} \iff \bar{C} = \arg \min_{C \in D} \tilde{d}_s(q, C), \quad (2)$$

where d_s is the distance between the query model \mathcal{Q} and the class C with respect to the spectrum sub-sequence s . Finally, the *average classification error* $\varepsilon \in [0, 1]$ is computed as the number of wrongly classified queries divided by the number of models.

Our analysis has been performed on the SHREC 2007 watertight data set (Fig. 2), which is grouped into 20 classes and contains 300 input shapes and 100 queries. To investigate the effectiveness of the selected information, each model has been re-meshed to 10K vertices. Different discretisations of the same or almost isometrically-deformed shapes are included to test the robustness of both the Laplacian spectrum and the eigenvalue selection. The results in the following subsections concern the selection of the eigenvalues on the training set in Fig. 2(a); the results in Section 4 are obtained by exploiting the selected features to query the training set with the complementary query set in Fig. 2(b).

3.1. Classification based on the first k eigenvalues

Since there are no theoretical results on the best suited value of k , a simple classification scheme based on the equation (1) and (2) has been investigated by using the sub-sequence $s = (1, \dots, k) \subseteq (1, \dots, k_{\max})$, where k varies from 1 the maximum number k_{\max} of computed eigenvalues. For the tests proposed in this section, we have chosen $k_{\max} := 500$. The distance d_s between the query \mathcal{Q} and the model \mathcal{M} is the L_2 -norm between their first k eigenvalues. Even if the un-normalised FEM spectrum is unbounded, our experiments have shown that the use of other distances (e.g., χ^2) do not sensibly improve the classification results. Finally, the proposed scheme selects the spectrum sub-sequence $[1, k_0]$, $k_0 \leq k$, corresponding to the sub-optimal minimization of the classification error.

Fig. 3 and Table 1 summarise the classification error computed on the normalised/un-normalised spectra and based on both the FEM and the cotangent discretisation. Note that the best results are produced by sequences of eigenvalues with small magnitudes and that the classification performance strongly decreases while increasing the magnitude of the eigenvalues. The growth of the classification error is effectively reduced by using re-meshed models and normalising the computed eigenvalues with their maximum. More precisely, the un-normalised (resp., normalised) FEM eigenvalues provides a classification error lower than 0.4 (resp., 0.3) with $8 \leq k \leq 13$ (resp., $7 \leq k \leq 19$) and the minimum $\varepsilon = 0.36$ (resp., $\varepsilon = 0.27$) is obtained with the first 12 (resp., 13) eigenvalues. For the re-meshed models, the un-normalised FEM spectrum provides a classification error lower than 0.4 by varying k between 7 and 13. The minimum error $\varepsilon = 0.36$ is achieved by selecting the first 11 eigenvalues. For the normalised FEM eigenvalues, a classification error lower than 0.25 is obtained with $13 \leq k \leq 18$ and for the minimum $\varepsilon = 0.22$ we use 15 eigenvalues.

Finally, for both the original and the remeshed models




Figure 2: (a) Data set and (b) query set.

we have computed the classification error using the first k cotangent eigenvalues. Considering the original data set, a classification error lower than 0.4 is obtained by selecting k between 12 and 18. For the re-meshed data set, values of ϵ lower than 0.4 are obtained for indices between 9 and 344. Concerning the original models, the best classification result $\epsilon = 0.37$ uses the first 13 eigenvalues, while for the re-meshed models the minimum error $\epsilon = 0.32$ is obtained with the first 59 eigenvalues.

3.2. Classification based on the Hill Climbing algorithm

The classification with respect to the *Hill Climbing with Backward* technique (HCBw, for short) selects a subsequence s of the shape spectrum that produces a sub-optimal minimisation of the classification error ϵ measured as in Equation (1). Assuming that s is initially empty, the HCBw algorithm iteratively adds to s those Laplacian eigenvalues that reduce ϵ . If ϵ stops to decrease, then the backward routine removes from s those elements that strongly reduce ϵ




Figure 3: Classification error (y-axis) on the first k eigenvalues, with k from 1 to 500 (x-axis).

		un-norm.		norm.	
		ϵ	k_0	ϵ	k_0
original models	FEM	0.30	11	0.27	13
	cot	0.37	13	-	-
re-meshed models	FEM	0.29	9	0.22	15
	cot	0.32	59	-	-

Table 1: Classification error ϵ computed by using the first k eigenvalues, $1 \leq k \leq 500$. Here, k_0 is the number of the eigenvalues that produces the minimum classification error.

(if any). If the backward routine cannot be iterated, then the HCBw algorithm tries to add new elements to s until possible. As shown in Table 2, the best HCBw classification result is obtained on the re-meshed data set using the FEM discretisation. Even if the best classification error is the same as the one shown in Table 1, the HCBw algorithm selects a smaller number of eigenvalues, which are among the first 50.

3.3. Classification based on the strong classifier produced by AdaBoost

The classification schema discussed in Sections 3.1 and 3.2 select a sequence of eigenvalues that minimise the classification error computed on the whole data set. In this section, *AdaBoost* [FS95] is applied to select the eigenvalues that minimise the classification error by considering each class of the data set separately. To this end, each class is associated to a set of eigenvalues such that they maximise the inter-class similarity and the dissimilarity among classes.

Using a set of positive and negative examples together with a large set of features, the goal of AdaBoost is to make the margin among the training examples as large as possible. As output, it produces a classification function based on a subset of features selected from the input set, which maximises the margin between the positive and negative examples. Given an input shape \mathcal{M} , the selection of the relevant features uses a set of binary classifiers (each of them

		un-norm.		norm.	
		ϵ	sel	ϵ	sel
original models	FEM	0.40	8	0.28	5
	cot	0.34	17	-	-
re-meshed models	FEM	0.37	13	0.22	10
	cot	0.34	7	-	-

Table 2: HCBw classification error ϵ with $1 \leq k \leq 500$ eigenvalues. The parameter sel corresponds to the number of selected eigenvalues.

associated to a single feature) that returns 1 if \mathcal{M} is recognised as belonging to the set of positive examples and -1 otherwise. Besides AdaBoost, other effective approaches for feature selection are based on the *Support Vector Machines* (SVM, for short) [BGV92, GGNZ06]. The main difference between SVM and AdaBoost is that the former relies on the definition of the most appropriate kernel to maximise the margin, while the latter achieves analogous results by using a fast greedy algorithm based weak classifiers. Moreover, the SVM requires to solve a quadratic programming problem, while the AdaBoost algorithm is based on linear programming. This reasons makes AdaBoost suitable for the efficient classification of high dimensional data.

In case of shape classification based on spectral information, each weak classifier is defined on a single eigenvalue of the spectrum. Indeed, AdaBoost is able to select the most relevant eigenvalues that maximise the margin between positive and negative examples. Moreover, each class is used in turn as positive example and the remaining classes are negative examples. In this way, a set of relevant eigenvalues has been obtained for each class. Among the shape features shared by similar models, we consider these sets of eigenvalues as those class descriptions that maximise the distance among models of different classes.

In our experiments, the weak classifier h_k classifies the shape model \mathcal{M} by using the k -th eigenvalues of its spectrum and it is defined as

$$h_k(\mathcal{M}) = \begin{cases} 1 & \max_{\mathcal{R} \in E^+} d_k(\mathcal{M}, \mathcal{R}) \leq th, \\ -1 & \text{otherwise,} \end{cases}$$

where E^+ is the set of positive examples, $d_k(\mathcal{M}, \mathcal{R})$ is the absolute value of the difference between the k -th eigenvalue of the models \mathcal{M} and \mathcal{R} , and $th = \max_{\mathcal{R}, \mathcal{Q} \in E^+} d_k(\mathcal{R}, \mathcal{Q})$, is a real number that is associated to the set of positive examples and represents the maximum distance between \mathcal{R} and \mathcal{Q} .

The following is a description of the AdaBoost algorithm used to produce the results summarised in Table 3.

- **Input examples:** $(\mathcal{M}_1, y_1), \dots, (\mathcal{M}_m, y_m)$, where $\mathcal{M}_i \in D$ and $y_i \in Y = \{+1, -1\}$, $i = 1, \dots, m$.

- **Initialisation:**

$$w_{0,i} := \frac{1}{2|E^+|} \text{ if } \mathcal{M} \in E^+, \quad w_{0,i} := \frac{1}{2|E^-|} \text{ otherwise,}$$

		un-norm.		norm.	
		ϵ	sel	ϵ	sel
original models	FEM	0.63	6.8	0.14	27.15
	cot	0.5	8.35	-	-
re-meshed models	FEM	0.68	7.3	0.03	30.6
	cot	0.05	32.1	-	-

Table 3: Classification error ϵ computed on the original/re-meshed data set with/without normalisation. sel is the average number of eigenvalues selected by AdaBoost.

where $|E^+|$ and $|E^-|$ are the number of positive and negative examples, respectively.

- **Iteration:** for $t = 1, \dots, T$,
 - train the weak classifiers by using the weights $w_{t,i}$;
 - select the weak classifier h_t producing the lowest classification error ϵ_t .
- **Update of the weights:** $w_{t+1,i} = \frac{1}{Z_t} w_{t,i} e^{-\alpha_t h_t(\mathcal{M}) y_i}$, where $\alpha_t = \frac{1}{2} \log \frac{1-\epsilon_t}{\epsilon_t}$ and Z_t is a normalisation factor such that $w_{t,i}$ ranges in $[0, 1]$.
- **Strong classifier:** $S = \sum_{t=1}^T \alpha_t h_t$.

The algorithm takes as input the set of positive and negative examples, $(\mathcal{M}_1, y_1), \dots, (\mathcal{M}_m, y_m)$, where \mathcal{M}_i is a model of the data set shown in Fig. 2(a) and $y_i \in \{+1, -1\}$ is the label representing a positive and negative example, respectively. AdaBoost iteratively trains the weak classifiers associated to the eigenvalues repeatedly in T iterations. During these iterations, the algorithm maintains a set of weights over the positive and negative examples. In particular, $w_{t,i}$ represents the weight of the example x_i at the iteration t . Once the weights have been initialised at the first step, they are iteratively updated on the basis of the incorrectly classified examples. In particular, at each iteration the weak classifier that produces the minimum classification error is selected and the error is used to update the weights of the input examples. Then, the weights of misclassified examples are increased and the weights of the correctly classified examples are reduced. This strategy forces the algorithm to focus only on hard examples.

The results summarised in Table 3 are obtained by running AdaBoost on the first 500 eigenvalues. The classification error has been computed on the original and re-meshed data set, using the two different Laplacian discretizations. Even though a larger number of iterations strongly reduces the classification error, we keep them as smaller as possible to reduce the overfitting on the data set. For this reason, the algorithm has been run with 20 iterations. The best results are obtained with the re-meshed models with both the normalised FEM and cotangent-based weights. Since Table 3 shows that the average number sel of eigenvalues selected by the algorithm is about the 3% of the cardinality of the input spectrum, we conclude that very few eigenvalues are sufficient to discriminate among the shape classes.

Fig. 4 shows that the selected eigenvalues for the classes




Figure 4: Selected eigenvalues for the class of (a) humans, (b) humanoids, (c) birds, and in (d) all the models in (a,b,c). The y-axis represents the eigenvalues, the x-axis their indexes; cross, circles and squares signs refer to the eigenvalues of human, humanoids and birds models, respectively.

of humans, humanoids, and birds are distributed on the computed spectrum. Even if humans and humanoids have similar shapes, the selected eigenvalues are distinct. Indeed, the eigenvalues with small magnitude (i.e., representing the overall shape) are overlapped and the eigenvalues with larger magnitude (i.e., representing local details) are separated. On the contrary, the selected eigenvalues for the class of the birds seldom overlap the class of humans and humanoids; in fact, AdaBoost selects those eigenvalues that strongly discriminate among different classes. Finally, Fig. 5 shows which intervals of the computed spectrum are induced by the selection of the relevant eigenvalues. Each class is represented by a column whose colored parts are the eigenvalues of the class members. It is interesting to notice that, for each class, a very small part of the spectrum can be used to distinguish among the dataset classes.

4. Shape retrieval based on the shape spectrum

Shape Retrieval is performed by computing the distance among the Laplacian spectra of the involved 3D models. This task can be achieved by selecting (a) the eigenvalues that maximise the retrieval performance or (b) the eigenvalues capable to identify the set of classes the query model most probably belongs to (classification task) and then by retrieving the models similar to the given query among the members of those classes. In the following, both the approaches have been investigated. Given a target eigenvalue λ , we can also compute a number of eigenvalues close




Figure 5: Selection of the relevant eigenvalues; the classes are plotted on the x-axis and the corresponding eigenvalues are reported on the y-axis.




Figure 6: Mean ADR (y-axis) for the data set in Fig. 2(a) and related to the first k , $1 \leq k \leq 500$, eigenvalues (x-axis).

to λ by applying a *spectral shift*; in fact, if λ is an eigenvalue of L , then $(\lambda - \sigma)$ is an eigenvalue of $(L - \sigma I)$.

The *Average Dynamic Recall* (ADR, for short) and the precision-recall diagram have been computed to evaluate the performance of the retrieval task [VtH07]. Given a query model Q , the $ADR \in [0, 1]$ represents how many models relevant to Q are retrieved among the Q shape models relevant to Q ; i.e., $ADR = \frac{1}{Q} \sum_{i=1}^Q \frac{r_i}{i}$, where r_i is the number of models relevant to Q and belonging to the first i retrieved models. Fig. 6 and Table 4 show the values of the ADR computed on the data set in Fig. 2(a). Each element has been used as query against the other elements of the data set and the value of the ADR has been averaged over all the queries (mean ADR). In Fig. 6, the ADR value has been computed by selecting the first k eigenvalues and by varying k between one and the cardinality of the computed shape spectrum. These diagrams show that the eigenvalues of smallest magnitude




Figure 7: Precision-recall diagrams related to Table 5.

		First k			
		un-norm.		norm.	
		mADR	k ₀	mADR	k ₀
original models	FEM	0.55	10	0.61	10
	cot	0.50	12	-	-
re-meshed models	FEM	0.56	10	0.62	15
	cot	0.57	12	-	-
		HCBw			
		un-norm.		norm.	
		mADR	sel	mADR	sel
original models	FEM	0.56	8	0.65	16
	cot	0.54	17	-	-
re-meshed models	FEM	0.56	8	0.66	15
	cot	0.57	6	-	-

Table 4: Mean ADR for the data set in Fig. 2(a).

(i.e., lower values of k) provide good performance due to their capability of coding the most representative shape features. In particular, the normalised FEM spectrum performs better than un-normalised ones and the cotangent eigenvalues. Note that shape comparison with a low (i.e., less than 5) or too high number of eigenvalues (i.e. more than 60) provides poor ADR values and compromises the retrieval performance. Table 4 summarise the results that maximise the mean ADR and reports also the eigenvalues that maximise the mean ADR selected through the HCBw algorithm. As for the classification, few selected eigenvalues yield good performance; in particular, the HCBw algorithm produces results that are better than the first k scheme and selects a smaller number of eigenvalues.

Table 5 reports the mean ADR that is obtained by considering the eigenvalues selected for the retrieval (see Table 4) and by combining classification and retrieval. To this end, the queries in Fig. 2(b) have been compared to the models in Fig. 2(a) by using the normalised FEM spectrum of the remeshed models. The three schemes based on the first k eigenvalues (first- k), the *Strong Classifier* (SC, for short), and the *Hill Climbing with Backward* (HCBw, for short) have been used to rank the classes for retrieval, while the

Classification	Retrieval	Classes	Mean ADR
SC	first- k	5	0.584
SC	HCBw	5	0.580
first- k	HCBw	5	0.571
HCBw	first- k	5	0.565
SC	first- k	10	0.563
SC	HCBw	10	0.560
HCBw	HCBw	5	0.560
HCBw	HCBw	10	0.555
HCBw	HCBw	15	0.555
first- k	HCBw	10	0.555
first- k	HCBw	15	0.555
-	HCBw	-	0.555
SC	HCBw	15	0.546
SC	first- k	15	0.542
HCBw	first- k	10	0.540
first- k	first- k	5	0.532
HCBw	first- k	15	0.530
first- k	first- k	10	0.530
first- k	first- k	15	0.530
-	first- k	-	0.530
-	random	-	0.337

Table 5: Mean Average Dynamic Recall for several classification schemes. The symbol – means that the retrieval has been performed without classification.

first k and the HCBw have been used directly for shape comparison. Then, the selected eigenvalues shown in the Tables 1, 2, 3 and the first 5, 10, 15 ranked classes have been exploited for classification and retrieval. For the strong classifier SC, the shape classes have been ranked according to the following probability $p(C|\mathcal{M}) = \frac{e^{SC(\mathcal{M})}}{e^{SC(\mathcal{M})} + e^{-SC(\mathcal{M})}}$ [LN07] while for the first k and the HCBw eigenvalues, the shape classes have been ranked using the distance d_s .

The eigenvalues selected to maximise the classification performance do not ensure good performance for retrieval. As previously described, the comparison among the query model and the members of the ranked classes is performed by using the eigenvalues that maximise the mean ADR on the data set as shown in Fig. 2(a). The results reported in Table 5 confirm that the retrieval performance are sensibly increased by combining both the classification and the retrieval tasks, in particular the results obtained by using the first k eigenvalues are over-performed by the schemes that use the eigenvalues selected by AdaBoost or by the HCBw algorithm. Fig. 7 shows the precision-recall diagrams corresponding to the best results of Table 5 for both retrieval with or without classification pre-processing.

Note that lower values of precision corresponding to low values of recall are explained as a higher number of false positive results. On the contrary, higher values of precision, corresponding to high values of recall, mean a lower number of false negative results. This remark highlights that even if almost all the schema have a similar behaviour with respect to false positives, the classification based on the HCBw combined with the first k eigenvalues for the shape comparison produces the lower number of false positives results. The strong classifier combined with the HCBw for shape com-

parison gives the lower number of false negatives results. Finally, the retrieval based on the randomly sampled spectrum and the retrieval obtained without the classification pre-processing task produces the worst performance.

5. Conclusions and future work

This paper shows that it is possible to select a set of Laplacian eigenvalues to strongly improve the classification and comparison results. In this context, the retrieval performance can be considerably improved by comparing the query shapes to a small subset of shape classes among the first ones ranked by a classification pre-processing. As future work, we will mainly investigate the localisation of those bands of eigenvalues that are shared by the models of the same class and the generalisation of the proposed approach to polygonal soups and point-sampled surfaces.

Acknowledgments Special thanks are given to the anonymous reviewers for their comments. This work has been partially supported by the FOCUS K3D EU CA.

References

- [BFF*07] BIASOTTI S., FALCIDIENO B., FROSINI P., GIORGI D., LANDI C., MARINI S., PATANÉ G., SPAGNUOLO M.: 3D shape description and matching based on properties of real functions. In *Eurographics-Tutorial* (2007), pp. 949–998.
- [BGV92] BOSER B. E., GUYON I. M., VAPNIK V. N.: A training algorithm for optimal margin classifiers. In *Proc. of the ACM Workshop on Computational Learning Theory* (1992), pp. 144–152.
- [BKS*05] BUSTOS B., KEIM D. A., SAUPE D., SCHRECK T., VRANIĆ D. V.: Feature-based similarity search in 3D object databases. *ACM Computing Surveys* 37, 4 (2005), 345–387.
- [CTSO03] CHEN D.-Y., TIAN X.-P., SHEN Y., OUHYOUNG M.: On visual similarity based 3D model retrieval. *Computer Graphics Forum* (2003), 223–232.
- [DZMC07] DYER R., ZHANG R. H., MOELLER T., CLEMENTS A.: An investigation of the spectral robustness of mesh laplacians. *Technical Report* (2007).
- [FS95] FREUND Y., SCHAPIRE R. E.: A decision-theoretic generalization of on-line learning and an application to boosting. In *Proc. of Computational Learning Theory* (1995).
- [FS99] FREUND Y., SCHAPIRE R. E.: A short introduction to boosting. In *In Proc. of the Sixteenth International Joint Conference on Artificial Intelligence* (1999), pp. 1401–1406.
- [GGNZ06] GUYON I., GUNN S., NIKRAVESH M., ZADEH L. (Eds.): *Feature Extraction, Foundations and Applications*, vol. 207 of *Studies in Fuzziness and Soft Computing*. Springer, 2006.
- [HLR05] HOU S., LOU K., RAMANI K.: Svm-based semantic clustering and retrieval of a 3d model database. In *Journal of Computer Aided Design and Application* (2005), pp. 155–164.
- [HR06] HOU S., RAMANI K.: A probability-based unified 3D shape search. In *Conference on Semantic and Digital Media Technologies* (2006), Lecture Notes in Computer Science, pp. 124–137.
- [JZ07] JAIN V., ZHANG H.: A spectral approach to shape-based retrieval of articulated 3D models. *Computer Aided Design* 39 (2007), 398–407.
- [LN07] LAGA H., NAKAJIMA M.: A boosting approach to content-based 3D model retrieval. In *Proc. of Computer Graphics and Interactive Techniques* (2007), pp. 227–234.
- [LN08] LAGA H., NAKAJIMA M.: Supervised learning of salient 2d views of 3D models. *The Journal of the Society for Art and Science* 7, 4 (2008), 124–131.
- [MÓ9] MÉMOLI F.: Spectral Gromov-Wasserstein distances for shape matching. In *Workshop on Non-Rigid Shape Analysis and Deformable Image Alignment* (2009).
- [MCHB07] MATEUS D., CUZZOLIN F., HORAUD R., BOYER E.: Articulated shape matching using locally linear embedding and orthogonal alignment. *IEEE International Conference on Computer Vision* (2007), 1–8.
- [MSF07] MARINI S., SPAGNUOLO M., FALCIDIENO B.: Structural shape prototypes for the automatic classification of 3D objects. *IEEE Computer Graphics and Applications* 27, 4 (2007), 28–37.
- [OFCD02] OSADA R., FUNKHOUSER T., CHAZELLE B., DOBKIN D.: Shape distributions. *ACM Transactions on Graphics* 21, 4 (2002), 807–832.
- [OK06] OHBUCHI R., KOBAYASHI J.: Unsupervised learning from a corpus for shape-based 3d model retrieval. In *Proc. of the Workshop on Multimedia Information Retrieval* (2006), pp. 163–172.
- [RBG*09] REUTER M., BIASOTTI S., GIORGI D., PATANÉ G., SPAGNUOLO M.: Discrete laplace-beltrami operators for shape analysis and segmentation. *Computers & Graphics* 33 (2009), 381–390.
- [Rus07] RUSTAMOV R. M.: Laplace-Beltrami eigenfunctions for deformation invariant shape representation. In *Proc. of the Symposium on Geometry processing* (2007), pp. 225–233.
- [RWP06] REUTER M., WOLTER F.-E., PEINECKE N.: Laplace-Beltrami spectra as Shape-DNA of surfaces and solids. *Computer-Aided Design* 38, 4 (2006), 342–366.
- [SF06] SHILANE P., FUNKHOUSER T.: Selecting distinctive 3D shape descriptors for similarity retrieval. In *Proc. of Shape Modeling and Applications* (2006), p. 18.
- [SOG09] SUN J., OVSJANIKOV M., GUIBAS L. J.: A concise and provably informative multi-scale signature based on heat diffusion. *Computer Graphics Forum* 28, 5 (2009), 1383–1392.
- [TV04] TIEU K., VIOLA P.: Boosting image retrieval. *International Journal of Computer Vision* 56, 1-2 (2004), 17–36.
- [TV08] TANGELDER J. W., VELTKAMP R. C.: A survey of content based 3D shape retrieval methods. *Multimedia Tools and Applications* 39, 3 (2008), 441–471.
- [VH07] VELTKAMP R. C., TER HAAR F.: *SHREC2007: 3D Shape Retrieval Contest*. Tech. rep., Utrecht University, 2007.
- [Wan09] WANG Y.: Approximating gradients for meshes and point clouds via diffusion metric. *Computer Graphics Forum* 28 (2009), 1497–1508(12).
- [ZvKD07] ZHANG H., VAN KAICK O., DYER R.: Spectral methods for mesh processing and analysis. In *Eurographics 2007 State-of-the-art Report* (2007), pp. 1–22.