



ELSEVIER

Contents lists available at ScienceDirect

Pattern Recognition Letters

journal homepage: www.elsevier.com/locate/patrec

Deep learning with geodesic moments for 3D shape classification

Lorenzo Luciano, A. Ben Hamza*

Concordia University, CIISE, 1455 de Maisonneuve Blvd. West, EV7.631, Montreal, QC, Canada

ARTICLE INFO

Article history:

Available online xxx

MSC:

41A05

41A10

65D05

65D17

Keywords:

Geodesic moments

Deep learning

Laplace–Beltrami

Stacked autoencoders

Shape classification

ABSTRACT

In this paper, we present a deep learning framework for efficient 3D shape classification using geodesic moments. Our approach inherits many useful properties from the geodesic distance, most notably the capture of the intrinsic geometric structure of 3D shapes and the invariance to isometric deformations. Moreover, we show the similarity between the convergent series of the geodesic moments and the inverse-square eigenvalues of the Laplace–Beltrami operator in the continuous setting. The proposed algorithm uses a two-layer stacked sparse autoencoder to learn deep features from geodesic moments by training the hidden layers individually in an unsupervised fashion, followed by a softmax classifier. Experimental results on three standard 3D shape benchmarks demonstrate superior performance of the proposed approach compared to existing methods.

© 2017 Elsevier B.V. All rights reserved.

1. Introduction

The availability of large 3D shape benchmarks has sparked a flurry of research activity in the development of efficient approaches for nonrigid shape analysis, including clustering, classification and retrieval [1–5]. Shape classification is a well-researched and fundamental problem in various domains, including pattern recognition, computer vision, and geometry processing. It basically refers to the process of organizing a database of shapes into a known number of classes, and the task is to assign new shapes to one of these classes.

Much of the recent work in 3D shape classification uses spectral shape descriptors, which represent a shape via a concise and compact signature aimed at facilitating the classification task. These shape representations are the building blocks of many shape classification algorithms, and may be broadly categorized into local and global descriptors. Local descriptors are usually defined on each point of the shape, while global descriptors are defined on the entire 3D shape. The category of Local descriptors include the global point signature include (GPS) [6], heat kernel signature (HKS) [7], wave kernel signature (WKS) [8], and spectral graph wavelet signature (SGWS) [9]. On the other hand, many global descriptors can be obtained from point signatures by integrating over the entire shape. One of the simplest global descriptors is Shape-DNA [10], which is defined as a truncated sequence of the eigenvalues of the

Laplace–Beltrami (LBO) arranged in increasing order of magnitude. Chaudhari et al. [11] introduced a new version of the GPS signature by setting the LBO eigenfunctions to unity. Gao et al. [12] developed a variant of Shape-DNA, referred to as compact Shape-DNA (cShape-DNA), which is an isometry-invariant signature resulting from applying the discrete Fourier transform to the area-normalized eigenvalues of the LBO. A comprehensive list of spectral descriptors can be found in [13,14].

More recently, deep learning has seen a rapid growth in popularity due largely to its great success in a variety of applications, including speech recognition, natural language processing, and image classification [15,16]. Inspired by information-processing in human nervous systems, deep learning extracts high-level features from data using multilayered neural networks. The success of deep neural networks has been greatly accelerated by using graphics processing units (GPUs), which have become the platform of choice for training large, complex learning systems [17–20]. Deep learning models such as convolutional neural networks, deep belief networks, and stacked autoencoders have recently been applied to 3D shape analysis to learn high-level features from 3D shapes. Bu et al. [17] introduced a deep belief networks based approach for 3D shape classification using a shape descriptor represented by a full matrix defined in terms of the geodesic distance and eigenfunctions of the LBO. Su et al. [18] proposed a 3D shape classification framework using multi-view convolutional neural networks by combining information from multiple 2D rendered images of a 3D shape into a single descriptor. Qi et al. [19] introduced a multiresolution filtering strategy with the aim at improving the perfor-

* Corresponding author.

E-mail address: hamza@ciise.concordia.ca (A. Ben Hamza).

mance of multi-view convolutional neural networks on 3D shape classification.

In this paper, we propose a deep learning framework, called deep geodesic moment (DeepGM) classifier, for 3D shape classification using geodesic moments and stacked sparse autoencoders. The geodesic moments are feature vectors derived from the integral of the geodesic distance on a shape, while stacked sparse autoencoders are deep neural networks consisting of multiple layers of sparse autoencoders that attempt to enforce a constraint on the sparsity of the output from the hidden layer. The proposed DeepGM approach learns deep discriminative features via deep learning with geodesic moments. It harnesses the geometric information from 3D shapes and then uses unsupervised autoencoders to extract high-level features from the geodesic moments. These high-level features are then fed into a supervised stacked sparse autoencoder architecture to learn a classification model from a training dataset of 3D shapes. We show that our model incorporates geometric features from shapes with the aim of designing a highly discriminative shape descriptor that yields better classification accuracy compared to existing methods.

The contributions of this paper are twofold: (1) We present an integrated framework for 3D shape classification that extracts discriminative geometric shape descriptors with geodesic moments. (2) We propose a classification approach that harnesses the power of deep learning to generate high-level features, which are in turn used within a stacked sparse autoencoder architecture with two hidden layers in an effort to accurately classify shapes in a database. Our experimental results show superior performance of the proposed framework over existing classification methods on several 3D shape benchmarks.

The rest of this paper is organized as follows. In Section 2, we briefly overview the Laplace–Beltrami operator and stacked sparse autoencoders. In Section 3, we propose a deep learning approach with geodesic moments for 3D shape classification using a two-layer stacked sparse autoencoder, and we discuss its main algorithmic steps. Experimental results for 3D shape classification are presented in Section 4. Finally, we conclude in Section 5 and point out some future work directions.

2. Background

In this section, we succinctly review the Laplace–Beltrami operator and stacked sparse autoencoders.

2.1. Laplace–Beltrami operator

Given a compact Riemannian manifold \mathbb{M} , the space $L^2(\mathbb{M})$ of all smooth, square-integrable functions on \mathbb{M} is a Hilbert space endowed with inner product $\langle f_1, f_2 \rangle = \int_{\mathbb{M}} f_1(\mathbf{x}) f_2(\mathbf{x}) da(\mathbf{x})$, for all $f_1, f_2 \in L^2(\mathbb{M})$, where $da(\mathbf{x})$ (or simply dx) denotes the measure from the area element of a Riemannian metric on \mathbb{M} . Given a twice-differentiable, real-valued function $f: \mathbb{M} \rightarrow \mathbb{R}$, the Laplace–Beltrami operator is defined as $\Delta_{\mathbb{M}} f = -\text{div}(\nabla_{\mathbb{M}} f)$, where $\nabla_{\mathbb{M}} f$ is the intrinsic gradient vector field and div is the divergence operator [21]. The LBO is a linear, positive semi-definite operator acting on the space of real-valued functions defined on \mathbb{M} , and it is a generalization of the Laplace operator to non-Euclidean spaces.

2.2. Stacked sparse autoencoders

An autoencoder is a neural network that learns to reproduce its input as its output. It is an unsupervised learning algorithm that learns features from unlabeled data using backpropagation via stochastic gradient descent, and has typically an input layer representing the original data, one hidden layer and an output layer.

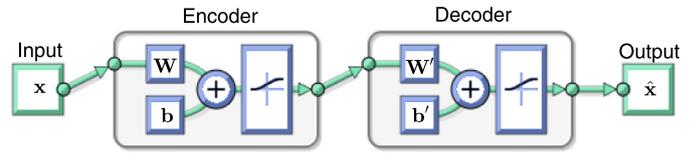


Fig. 1. Graphical diagram of an autoencoder.

An autoencoder is comprised of an encoder and a decoder, as illustrated in Fig. 1.

The encoder, denoted by f_{θ} , maps an input vector $\mathbf{x} \in \mathbb{R}^q$ to a hidden representation (referred to as code, activations or features) $\mathbf{a} \in \mathbb{R}^r$ via a deterministic mapping

$$\mathbf{a} = f_{\theta}(\mathbf{x}) = \sigma(\mathbf{W}\mathbf{x} + \mathbf{b}), \quad (1)$$

parameterized by $\theta = \{\mathbf{W}, \mathbf{b}\}$, where $\mathbf{W} \in \mathbb{R}^{r \times q}$ and $\mathbf{b} \in \mathbb{R}^r$ are the encoder weight matrix and bias vector, and σ is a nonlinear element-wise activation function such as the logistic sigmoid or hyperbolic tangent. The decoder, denoted by $g_{\theta'}$, maps back the hidden representation \mathbf{h} to a reconstruction $\hat{\mathbf{x}}$ of the original input \mathbf{x} via a reverse mapping

$$\hat{\mathbf{x}} = g_{\theta'}(\mathbf{a}) = \sigma(\mathbf{W}'\mathbf{a} + \mathbf{b}'), \quad (2)$$

parameterized by $\theta' = \{\mathbf{W}', \mathbf{b}'\}$, where $\mathbf{W}' \in \mathbb{R}^{q \times r}$ and $\mathbf{b}' \in \mathbb{R}^q$ are the decoder weight matrix and bias vector, respectively. The encoding and decoding weight matrices \mathbf{W} and \mathbf{W}' are usually constrained to be of the form $\mathbf{W}' = \mathbf{W}^T$, which are referred to as tied weights. Assuming the tied weights case for simplicity, the parameters $\{\mathbf{W}, \mathbf{b}, \mathbf{b}'\}$ of the network are often optimized by minimizing the squared error $\sum_{i=1}^N \|\mathbf{x}_i - \hat{\mathbf{x}}_i\|_2^2$, where N is the number of samples in the training set, \mathbf{x}_i is the i th input sample and $\hat{\mathbf{x}}_i$ is its reconstruction.

To penalize large weight coefficients in an effort to avoid overfitting the training data and also to encourage sparsity of the output from the hidden layer, the following objective function is minimized instead

$$\mathcal{L}(\mathbf{W}, \mathbf{b}, \mathbf{b}') = \frac{1}{2} \sum_{i=1}^N \|\mathbf{x}_i - \hat{\mathbf{x}}_i\|_2^2 + \frac{\lambda}{2} \|\mathbf{W}\|_F^2 + \beta \sum_{j=1}^N \text{KL}(\rho \| \hat{\rho}_j), \quad (3)$$

where λ is a regularization parameter that determines the relative importance of the sum-of-squares error term and the weight decay term, and β is the weight of the sparsity regularization term. This sparsity regularizer is the Kullback–Leibler divergence $\text{KL}(\rho \| \hat{\rho}_j)$, which is a dissimilarity measure between ρ and $\hat{\rho}_j$, and it is defined as

$$\text{KL}(\rho \| \hat{\rho}_j) = \rho \log \frac{\rho}{\hat{\rho}_j} + (1 - \rho) \log \frac{1 - \rho}{1 - \hat{\rho}_j}, \quad (4)$$

where $\hat{\rho}_j$ is the average activation value of the hidden unit j and ρ is its desired value which is typically small.

An autoencoder with multiple hidden layers is referred to as a stacked or deep autoencoder. A stacked sparse autoencoder is a deep neural network consisting of multiple layers of stacked encoders from several sparse autoencoders. This stacked network is pre-trained layer by layer in an unsupervised fashion, where the output from the encoder of the first autoencoder is the input of the second autoencoder, the output from the encoder of the second autoencoder is the input to the third autoencoder, and so on. After pre-training, the entire stacked sparse autoencoder can be trained using backpropagation to fine-tune all the parameters of the network. For classification purposes, the final hidden layer of the stacked sparse autoencoder is linked to the softmax layer. The latter is trained in a supervised manner using labels for the training data.

3. Method

In this section, we describe in detail the proposed deep learning approach to 3D shape classification. We start by defining the continuous geodesic moments on a Riemannian manifold. We show that the second geodesic moment reduces to the integral of the geodesic distance [22], while the third geodesic moment is related to the heat kernel [7]. Then, we build the geodesic moments in the discrete setting, and we provide the main algorithmic steps of our classification framework.

3.1. Geodesic moments

Let \mathbb{M} be a connected, compact Riemannian manifold, without boundary. For a fixed (source) point y in \mathbb{M} , we define the k th order geodesic central moment (or simply k th geodesic moment) about y as

$$\mu_k(y) = \int_{\mathbb{M}} d^k(x, y) \Delta_{\mathbb{M}} d(x, y) da(x), \quad (5)$$

where $d(x, y)$ is the geodesic distance starting from point y [23]. Note that the geodesic distance function $f(x) = d(x, y)$ is a solution to the Eikonal equation $\|\nabla_{\mathbb{M}} f\| = 1$ on \mathbb{M} , with the initial condition $f(y) = 0$. In other words, f is the geodesic distance function from a source point y to the remaining points on the manifold. The first two moments are

$$\mu_0 = \int_{\mathbb{M}} \Delta_{\mathbb{M}} f da = - \int_{\mathbb{M}} \operatorname{div}(\nabla_{\mathbb{M}} f) da = 0$$

and

$$\mu_1 = \int_{\mathbb{M}} f \Delta_{\mathbb{M}} f da = \int_{\mathbb{M}} \|\nabla_{\mathbb{M}} f\|^2 da = \int_{\mathbb{M}} da = a,$$

respectively, where $a = \operatorname{area}(\mathbb{M})$. Note that $\mu_0 = 0$ and μ_1 do not depend on the source point y .

Theorem 1. For all $k \in \mathbb{N}$, the k th geodesic moment is given by

$$\mu_k(y) = k \int_{\mathbb{M}} d^{k-1}(x, y) da(x), \quad \text{for all } y \in \mathbb{M}. \quad (6)$$

Proof. Denote by f the geodesic distance function. For all $k \in \mathbb{N}$, we have

$$\begin{aligned} \Delta_{\mathbb{M}} f^{k+1} &= -\operatorname{div}(\nabla_{\mathbb{M}} f^{k+1}) = -\operatorname{div}((k+1)f^k \nabla_{\mathbb{M}} f) \\ &= (k+1)f^k \Delta_{\mathbb{M}} f - \langle (k+1)\nabla_{\mathbb{M}} f^k, \nabla_{\mathbb{M}} f \rangle \\ &= (k+1)f^k \Delta_{\mathbb{M}} f - \langle k(k+1)f^{k-1} \nabla_{\mathbb{M}} f, \nabla_{\mathbb{M}} f \rangle \\ &= (k+1)f^k \Delta_{\mathbb{M}} f - k(k+1)f^{k-1} \langle \nabla_{\mathbb{M}} f, \nabla_{\mathbb{M}} f \rangle \\ &= (k+1)f^k \Delta_{\mathbb{M}} f - k(k+1)f^{k-1} \|\nabla_{\mathbb{M}} f\|^2 \\ &= -(k+1)f^{k-1} (-f \Delta_{\mathbb{M}} f + k \|\nabla_{\mathbb{M}} f\|^2) \\ &= (k+1)f^k \Delta_{\mathbb{M}} f - k(k+1)f^{k-1} \|\nabla_{\mathbb{M}} f\|^2 \\ &= (k+1)f^k \Delta_{\mathbb{M}} f - k(k+1)f^{k-1}. \end{aligned}$$

Integrating and applying the divergence theorem, we obtain

$$\mu_k(y) = k \int_{\mathbb{M}} f^{k-1}(x) da(x), \quad (7)$$

which concludes the proof. \square

It is worth pointing out that (\mathbb{M}, d) equipped with the geodesic distance d is a metric space, and the quantity $\int_{\mathbb{M}} d^{k-1}(x, y) da(x)$ is often referred to as the $(k-1)$ th central moment in metric spaces.

Using (6), it can readily be shown that the k th geodesic moments are positive and satisfy the following inequality

$$0 = \mu_0 < \mu_k(y) \leq k a \rho^{k-1}, \quad k \geq 1, \quad (8)$$

where $\rho = \sup_{x, y \in \mathbb{M}} d(x, y)$ is the diameter of \mathbb{M} . The mesh diameter is the longest distance between two vertices in the mesh, and

is invariant under global isometries. It is a measure of mesh compactness in the sense that it measures the dissimilarity between the most dissimilar pair of vertices in the mesh. If $\rho < 1$, then the series $\sum_{k \geq 1} \mu_k(y)$ is convergent for all $y \in \mathbb{M}$.

Since $\Delta_{\mathbb{M}}$ is a positive semi-definite operator, its eigenvalues are nonnegative and may be written in increasing order as $0 = \lambda_0 < \lambda_1 \leq \lambda_2 \leq \dots$ with $\lambda_k \nearrow \infty$ as $k \nearrow \infty$. Moreover, the series $\sum_{k \geq 1} 1/\lambda_k^2$ is convergent. Notice the similarity between the convergent series of the geodesic moments and inverse-square eigenvalues of the LBO.

For a fixed source point y , if $d(x, y) \geq 1$ for all x , then the geodesic moments are nonnegative and may be written in increasing order as $0 = \mu_0(y) < \mu_1(y) \leq \mu_2(y) \leq \dots$

Relation to integral of geodesic distance: The second geodesic moment is given by

$$\mu_2(y) = 2 \int_{\mathbb{M}} d(x, y) dx, \quad (9)$$

which is proportional to the aggregated geodesic distance from y to the remaining points on the manifold. The aggregated geodesic distance (also called global geodesic function) is invariant to isometric transformations, and has been used successfully in generating robust Reeb graphs for shape matching and classification [22]. A small value of $\mu_2(y)$ indicates that a distance from y to arbitrary points on \mathbb{M} is relatively small (i.e., the point y is closer to the center of the object).

Relation to heat kernel: Using Varadhan's formula [21], we may express the third geodesic moment as

$$\mu_3(y) = 3 \int_{\mathbb{M}} d^2(x, y) dx = \lim_{t \rightarrow 0} -12t \int_{\mathbb{M}} \log h_t(x, y) dx, \quad (10)$$

where $h_t(x, y)$ is called the heat kernel, which measures the amount of heat that is propagated or transferred from point x to point y in time t .

3.2. Discrete geodesic moments

Let \mathbb{M} be a 3D shape represented by a triangle mesh with m vertices. In the discrete setting, we may write the k th geodesic moment in (6) as

$$\mu_k(j) \approx k \sum_{i=1}^m d_{ij}^{k-1} a_i, \quad (11)$$

where a_i is the area of the Voronoi cell at vertex i , and d_{ij} is the geodesic distance between mesh vertices i and j . We refer to the p -dimensional vector $\boldsymbol{\mu}_j = (\mu_1(j), \dots, \mu_p(j))$ consisting of the first p moments as the *geodesic moment point signature* at vertex j . Hence, we may represent the shape \mathbb{M} by an $m \times p$ geodesic moment matrix $\mathbf{M} = (\boldsymbol{\mu}_1, \dots, \boldsymbol{\mu}_m)^T$ of m signatures, each of which is of length p . In other words, the rows of \mathbf{M} are data points and the columns are features.

As the number of mesh vertices differs from one shape to another in a dataset of shapes, the geodesic moments matrix may not be a convenient representation for shape analysis tasks such as retrieval and classification. To circumvent this limitation, we represent the shape \mathbb{M} by the $p \times p$ matrix $\mathbf{S} = \mathbf{M}^T \mathbf{M}$, which we refer to as the geodesic moment descriptor of the shape. This matrix may be regarded as an uncentered scatter matrix. In addition to being independent of the number of mesh vertices, the geodesic moment descriptor enjoys a number of desirable properties including simplicity, compactness, invariance to isometric deformations, and computational feasibility.

3.3. Proposed algorithm

The goal of 3D shape classification is to accurately predict the target class for each 3D shape in a dataset. This is typically done

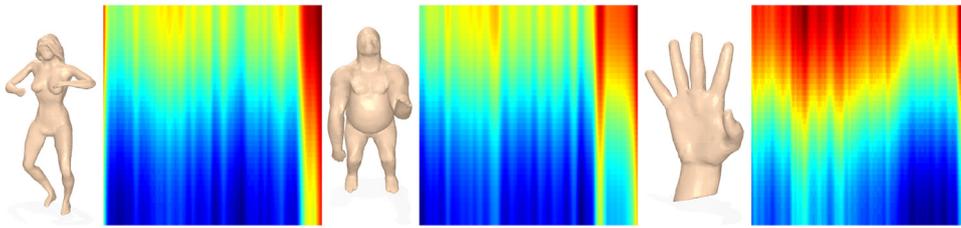


Fig. 2. Geodesic moment signatures for three shapes (woman, gorilla, and hand) from three different classes of the SHREC-2011 dataset.

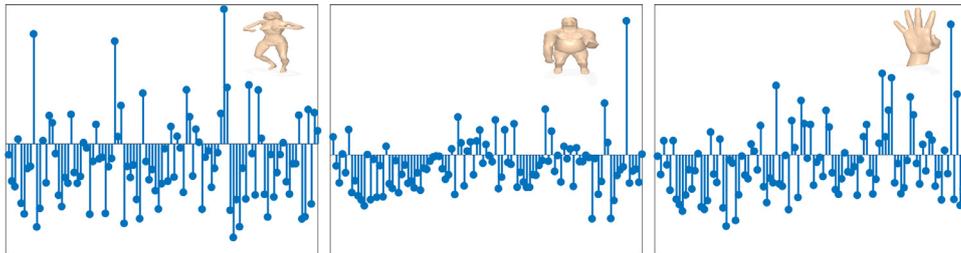


Fig. 3. DeepGM learned features for three shapes (woman, gorilla, and hand) from three different classes of the SHREC-2011 dataset.

by extracting discriminative features from 3D shapes, followed by using a learning technique to classify these shapes. The available data \mathcal{X} for shape classification is usually split into two disjoint subsets: the training set $\mathcal{X}_{\text{train}}$ for learning, and the test set $\mathcal{X}_{\text{test}}$ for testing. The training and test sets are customarily selected by randomly sampling a set of training instances from \mathcal{X} for learning and using the rest of instances for testing.

Our proposed DeepGM approach to 3D shape classification consists of two major steps. In the first step, we compute the $p \times p$ matrix $\mathbf{S}_i = \mathbf{M}_i^T \mathbf{M}_i$ for each shape \mathbb{M}_i in the dataset $\mathcal{D} = \{\mathbb{M}_1, \dots, \mathbb{M}_n\}$, where \mathbf{M}_i is the geodesic moment matrix and p is the number of geodesic moments. Then, each matrix \mathbf{S}_i is reshaped into a p^2 -dimensional feature vector \mathbf{x}_i by stacking its columns one underneath the other. Subsequently, all feature vectors \mathbf{x}_i of all n shapes in the dataset are arranged into a $n \times p^2$ data matrix $\mathbf{X} = (\mathbf{x}_1, \dots, \mathbf{x}_n)^T$. Fig. 2 displays the geodesic moment matrices of three shapes (woman, gorilla, and hand) from three different classes of the SHREC-2011 dataset [24].

In the second step, we use a 2-layer SSAE to learn deep features by training the hidden layers individually in an unsupervised manner. Then, we train a softmax layer to classify the deep features learned by the second autoencoder. Finally, the final network is trained (tuned) in a supervised fashion to determine the optimal parameters. Fig. 3 displays the step plots of DeepGM learned features of three shapes from three different classes of the SHREC-2011 dataset.

The task in multiclass classification is to assign a class label to each input example. More precisely, given a training data of the form $\mathcal{X}_{\text{train}} = \{(\mathbf{x}_i, y_i)\}$, where $\mathbf{x}_i \in \mathbb{R}^{p^2}$ is the i th example and $y_i \in \{1, \dots, K\}$ is its i th class label, we aim at finding a learning model that contains the optimized parameters from the SSAE algorithm. Then, the trained deep learning model is applied to a test data $\mathcal{X}_{\text{test}}$, resulting in predicted labels \hat{y}_i of new data. These predicted labels are subsequently compared to the labels of the test data to evaluate the classification accuracy of the model.

To assess the performance of the proposed framework, we employed two commonly used evaluation criteria, the confusion matrix and accuracy, which will be discussed in more detail in the next section. Algorithm 1 summarizes the main algorithm steps of our DeepGM approach.

Algorithm 1 DeepGM classifier.

Input: Dataset $\mathcal{D} = \{\mathbb{M}_1, \dots, \mathbb{M}_n\}$ of n shapes

- 1: **for** $i = 1$ to n **do**
- 2: Compute the $m \times p$ geodesic moment matrix \mathbf{M}_i for each 3D shape \mathbb{M}_i , where m is the number of vertices
- 3: Compute the $p \times p$ matrix $\mathbf{S}_i = \mathbf{M}_i^T \mathbf{M}_i$, and reshape it into a p^2 -dimensional vector \mathbf{x}_i
- 4: **end for**
- 5: Arrange all the feature vectors \mathbf{x}_i into a $n \times p^2$ data matrix $\mathbf{X} = (\mathbf{x}_1, \dots, \mathbf{x}_n)$
- 6: Apply a 2-layer stacked sparse autoencoder with a softmax classifier on \mathbf{X} to find the n -dimensional vector $\hat{\mathbf{y}}$ of predicted class labels.

Output: n -dimensional vector $\hat{\mathbf{y}}$ containing predicted class labels for each 3D shape

4. Experiments

In this section, we conduct extensive experiments to assess the performance of the proposed DeepGM approach in 3D shape classification. The effectiveness of our approach is validated by performing a comprehensive comparison with several shape classification methods.

Datasets: The effectiveness of the proposed framework is evaluated on three standard and publicly available 3D shape benchmarks [24–26]: SHREC-2010, SHREC-2011 and SHREC-2015. Sample shapes from these datasets are shown in Fig. 4.

Performance evaluation measures: The performance of a classifier is usually assessed by applying it to test data with known target values and comparing the predicted values with the known values. One important way of evaluating the performance of a classifier is to compute its confusion matrix, which is a $K \times K$ matrix that displays the number of correct and incorrect predictions made by the classifier compared with the actual classifications in the test set, where K is the number of classes.

Another intuitively appealing measure is the classification accuracy, which is a summary statistic that can be easily computed from the confusion matrix as the total number of correctly classified instances (i.e., diagonal elements of confusion matrix) divided by the total number of test instances. Alternatively, the accuracy of



Fig. 4. Sample shapes from SHREC-2010 (top), SHREC-2011 (middle), and SHREC-2015 (bottom).

a classification model on a test set may be defined as follows

$$\begin{aligned} \text{Accuracy} &= \frac{\text{Number of correct classifications}}{\text{Total number of test cases}} \\ &= \frac{|\mathbf{x} : \mathbf{x} \in \mathcal{X}_{\text{test}} \wedge \hat{y}(\mathbf{x}) = y(\mathbf{x})|}{|\mathbf{x} : \mathbf{x} \in \mathcal{X}_{\text{test}}|}, \end{aligned} \quad (12)$$

where $y(\mathbf{x})$ is the actual (true) label of \mathbf{x} , and $\hat{y}(\mathbf{x})$ is the label predicted by the classification algorithm. A correct classification means that the learned model predicts the same class as the original class of the test case. The error rate is equal to one minus accuracy.

Baseline methods: Using the aforementioned 3D shape benchmarks, we carry out a comprehensive comparison between the proposed DeepGM framework and several baseline methods, including Shape-DNA [10], compact Shape-DNA [12], GPS embedding [11], and F1-, F2-, and F3-features [27]. In order to show the power of deep learning in further improving the classification accuracy results, we also compared DeepGM to the shallow geodesic moment (GM) approach, which employs a one-vs-all SVM classifier (i.e., using SVM in Step 6 of Algorithm 1).

Implementation details: All experiments were conducted on a Dell Optiplex 9020 Windows desktop computer with an Intel Core i7 running 3.4 GHz and 32GB RAM; and all the algorithms were implemented in MATLAB. For training, we used the $p^2 - h_1 - h_2 - K$ architecture for a stacked sparse autoencoder with two layers as illustrated in Fig. 5, where p^2 is the size of the input, h_1 is the size of the hidden layer for the first autoencoder, h_2 is the size of the hidden layer for the second autoencoder, and K is the number of classes in the dataset. We used the logistic sigmoid function as an activation function for both autoencoders. We also set the regularization parameter to $\lambda = 0.0001$, the target sparsity probability to $\rho = 0.15$, and the weight of the sparsity regularization term to $\beta = 3$.

The length of a shape descriptor is usually domain specific and often chosen empirically through experimentation. As will be discussed later, our extensive experiments reveal that a number of geodesic moments between 20 and 30 gives better classification accuracy. In particular, we set the number of moments to 20 for SHREC-2010 and SHREC-2011, and to 30 for SHREC-2015. In other

Table 1

Classification accuracy results on the SHREC-2010 dataset. Boldface numbers indicate the best classification performance.

Method	Average accuracy %
F1-features	76.83 ± 2.77
F2-features	77.17 ± 4.01
F3-features	75.83 ± 3.95
Shape-DNA	82.67 ± 1.96
cShape-DNA	78.50 ± 5.58
GPS-embedding	87.17 ± 3.60
GM	92.17 ± 4.17
DeepGM	96.33 ± 1.05

words, each shape in the SHREC-2010 and SHREC-2011 datasets is represented by a 400-dimensional geodesic feature vector, while each shape in SHREC-2015 is described by a 900-dimensional feature vector.

The sizes of the hidden layers for the first and second autoencoders are set to $h_1 = 200$ and $h_2 = 100$, respectively. In our DeepGM approach, we use the features learned by the second autoencoder to perform shape classification. More specifically, we classify the 100-dimensional deep feature vectors by training a softmax layer in a supervised fashion using labels for the training data. In all our experiments, we selected a 70%/30% split between training and testing, respectively.

For the baseline methods shape-DNA, GPS embedding, F1-, F2-, and F3-features, we set the number of retained eigenvalues to 10. For cShape-DNA, the dimension of the signature was set to 33 as suggested in [12].

4.1. SHREC-2010 dataset

The SHREC-2010 dataset consists of 200 nonrigid 3D models, which are classified into 10 categories. The models are represented as watertight triangle meshes [25]. Each category contains 20 objects in a different posture. The 10 categories that make up the dataset are: ants, crabs, hands, humans, octopus, pliers, snakes, spectacle, spiders and teddy bear.

Performance evaluation: The SHREC-2010 dataset is randomly divided into 70% for training and 30% for testing, yielding a test set of 60 samples. We first train the model using the training data to learn the classification model. Then, the resulting model is used on the test data to derive classification results. The confusion matrix displayed in Fig. 6 shows that the proposed DeepGM approach was able to classify most of the objects correctly, except for one instance of the spider model which was misclassified as a crab, indicating that DeepGM is able to capture discriminative features to distinguish between various 3D shapes.

Results: In our DeepGM approach, each 3D shape in the SHREC-2010 dataset is represented using 20 geodesic moments, resulting in a data matrix \mathbf{X} of size 400×200 .

We compare the proposed DeepGM method to shape-DBA, compact shape-DNA, GPS embedding, F1-, F2-, F3-features, and the GM approach. For each method, we followed the standard practice of running the algorithm 10 times with different randomly selected training and testing datasets, and then we recorded the mean and standard deviation for each method. Table 1 displays the classification accuracy rates for all methods. As can be seen, DeepGM achieves better performance than Shape-DNA, cShape-DNA, GPS embedding, F1-, F2-, F3-features, and the GM classifier. The classification accuracy of DeepGM is 96.33%, which shows improvements of 4.16% and 9.16% over the best performing classifiers GM approach and GPS embedding, respectively. In addition to plotting the average classification accuracy rates, Fig. 7 shows the error bars of DeepGM and the baseline methods on the SHREC-2010 dataset.

Table 2

Classification accuracy results on the SHREC-2011 dataset. Boldface numbers indicate the best classification performance.

Method	Average accuracy %
F1-features	83.78 ± 3.03
F2-features	80.83 ± 2.24
F3-features	80.33 ± 2.15
Shape-DNA	85.39 ± 2.36
cShape-DNA	88.77 ± 1.77
GPS-embedding	83.22 ± 1.88
GM	94.11 ± 1.28
DeepGM	97.89 ± 0.57

Table 3

Classification accuracy results on the SHREC-2015 dataset. Boldface numbers indicate the best classification performance.

Method	Average accuracy %
F1-features	56.03 ± 4.46
F2-features	50.86 ± 3.54
F3-features	62.66 ± 1.90
Shape-DNA	61.17 ± 3.38
cShape-DNA	74.80 ± 1.41
GPS-embedding	63.40 ± 1.73
GM	83.34 ± 1.88
DeepGM	93.03 ± 0.64

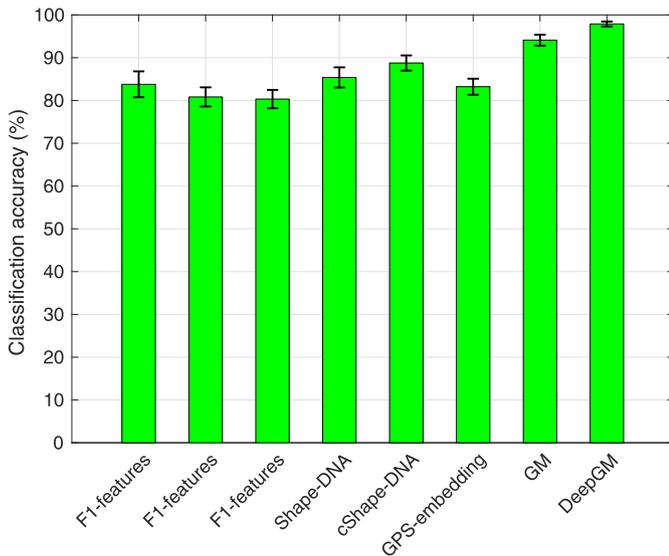


Fig. 9. Classification accuracy rates with error bars for DeepGM and baseline methods on SHREC-2011.

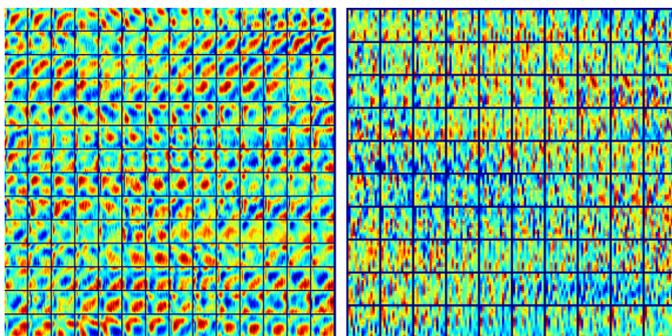


Fig. 10. DeepGM learned weights from the first layer (left) and second layer (right) on the SHREC-2011 dataset.

the GM approach which uses a one-vs-all multiclass SVM classifier. The average accuracy of DeepGM is 97.89% with performance improvements of 9.12% and 12.50% over cShape-DNA and Shape-DNA, respectively. DeepGM also yields a 3.78% accuracy improvement over the GM using classifier, indicating the advantage of using deep learning models over shallow ones.

The learned weights from the first and second autoencoders are shown in Fig. 10. The features learned by the second autoencoder are depicted in Fig. 11. As can be shown, the features are quite similar for shapes from the same class, meaning that DeepGM is robust to nonrigid deformations.

4.3. SHREC-2015 dataset

The SHREC-2015 dataset consists of 1200 3D watertight triangle meshes, which are classified into 50 categories. The models in each category are obtained by transforming the original 3D meshes of the same category.

Performance evaluation: We randomly selected 30% shapes in SHREC-2015 to hold out for the test set, and the remaining shapes for training. That is, the training and test sets consist of 840 and 160 shapes, respectively. First, we train a stacked sparse autoencoder with two hidden layers on the training data to learn the deep classification model. Then, we use the resulting, trained model on the test data to predict the class labels.

Results: Each 3D shape in the SHREC-2015 dataset is represented by a 900-dimensional feature vector, resulting in a data matrix \mathbf{X} of size 900×1200 . In a bid to obtain reliable results, we repeated the experiments 10 times for the proposed DeepGM approach as well as for each of the baseline methods, and we recorded the resulting classification accuracy rates. The average accuracy results along with the standard deviations are shown in Table 3. Fig. 12 displays the SHREC-2015 classification accuracy rate for each method along with the associated error. As can be seen, DeepGM outperforms all the baseline methods. The average accuracy of DeepGM is 93.03%, with performance improvements of 18.23% and 29.63% over the best performing baseline methods of cShape-DNA and GPS-embedding, respectively. DeepGM also yields a 9.69% accuracy improvement over the GM approach, which uses a one-vs-all multiclass SVM classifier. These substantial improvements in accuracy rates over the baseline methods dataset demonstrate the robustness of DeepGM in 3D shape classification.

4.4. Parameter sensitivity

We tested the performance of the proposed DeepGM approach using different values for the number of geodesic moments. Our experiments show that a number of geodesic moments in the range of 20–30 is usually sufficient to capture the discriminative features from 3D shapes for better classification accuracy, as depicted in Fig. 13 for all the three benchmarks used in experimentation.

5. Discussion

It is worth pointing out that an image descriptor, called patch geodesic moments, has been previously applied to 2.5D face recognition [28]. While the patch geodesic moments are computed via the local geodesic distance on range images only, the proposed geodesic moments are defined via spectral geometry and can be computed on both range data and 3D shapes. Moreover, we adopt a completely different strategy that uses the global geodesic distance in conjunction with the Laplace–Beltrami operator to represent the shape information. In other words, our approach differs

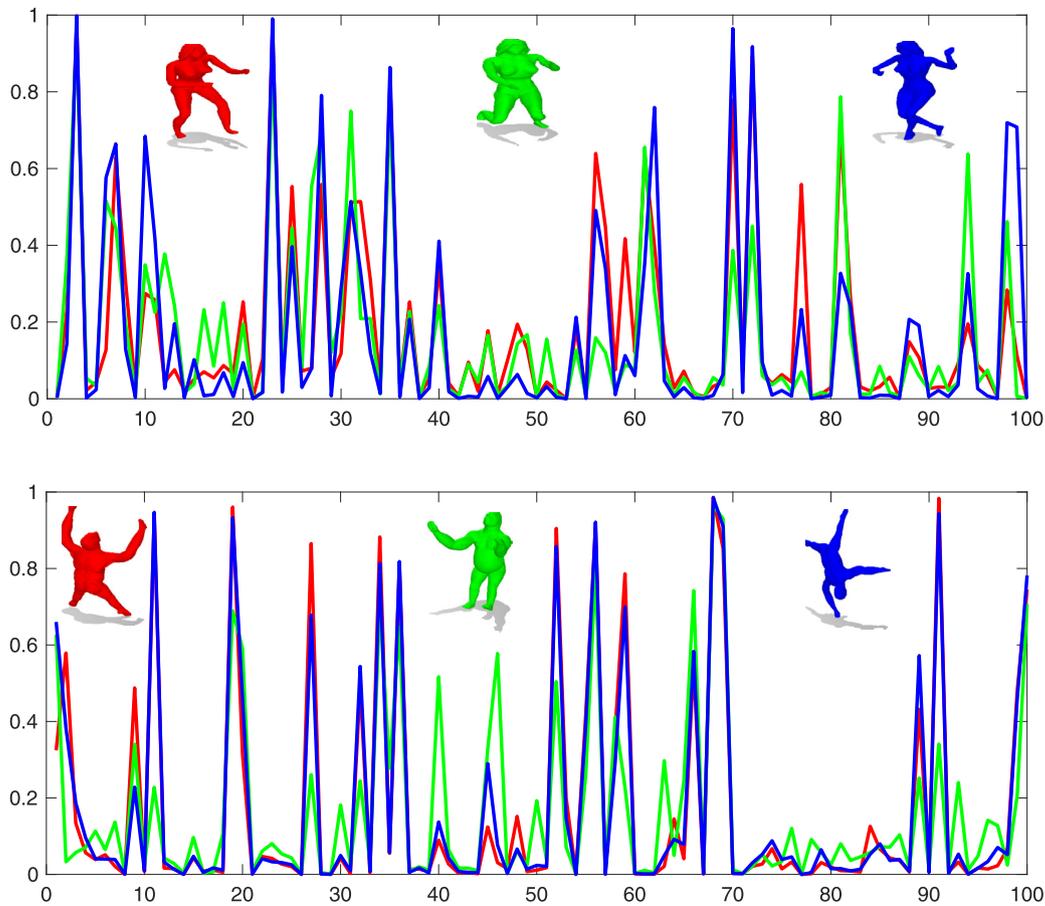


Fig. 11. DeepGM learned features for three women models (top) and three gorilla models (bottom) from the SHREC-2011 dataset.

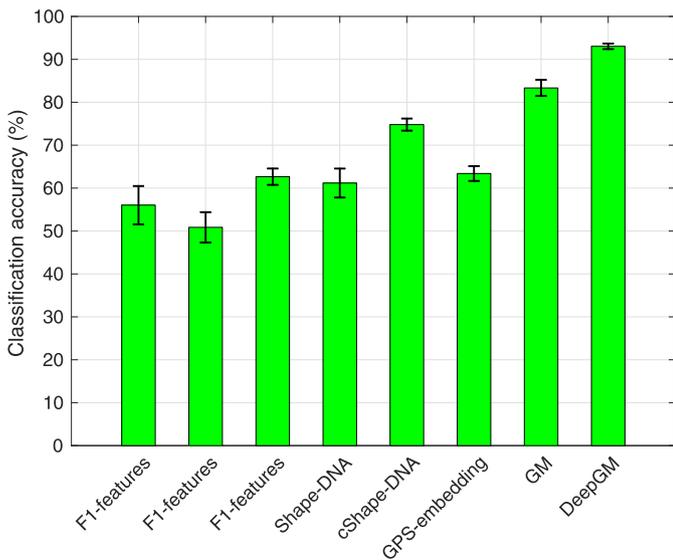


Fig. 12. Classification accuracy rates with error bars for DeepGM and baseline methods on SHREC-2015.

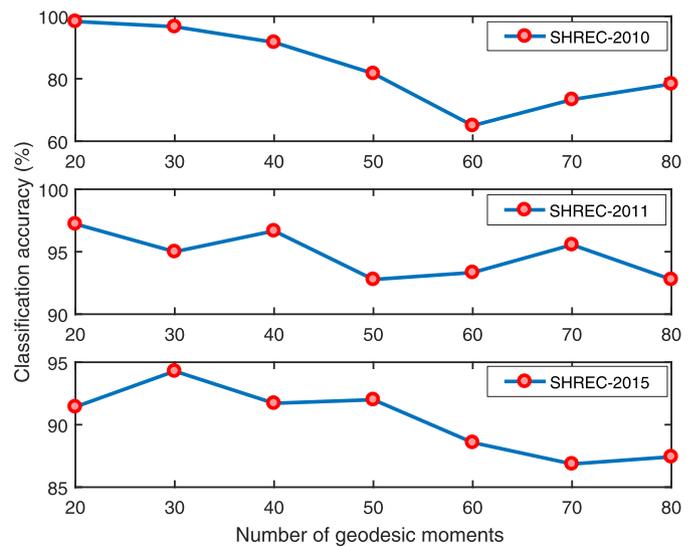


Fig. 13. Classification accuracy vs. number of geodesic moments.

in the way the geodesic moments are defined and computed. To enrich the geodesic moments based shape representation, we employed a 2-layer sparse stacked autoencoder to learn deep features by training the hidden layers individually in an unsupervised manner, resulting in a compact yet discriminative deep shape descriptor. The proposed DeepGM framework embodies several key ideas

that provide significant benefits over other existing shape classification methods. DeepGM represents the 3D shapes in a dataset by a low-dimensional feature vectors, which are arranged into a data matrix. The latter is then fed into a sparse stacked autoencoder to learn deep features, which are subsequently classified by training a softmax layer.

We performed extensive experiments to demonstrate that our DeepGM approach achieves state-of-the-art performance on different datasets. In particular, we included a shallow GM classifier in our experiments in order to show not only the efficiency of the geodesic model, but also the power of the deep learning in significantly improving the classification results. The experimental results show improved classification accuracy results of the DeepGM approach over the shallow GM classifier, with improvements of 4.16%, 3.78% and 9.69% on the SHREC-2010, SHREC-2011, and SHREC-2015 datasets, respectively.

A major issue with most machine learning models is overfitting, which needs to be adequately addressed. To address this concern, we resorted to a number of techniques that provide effective remedies to overfitting: (1) the use of the L_2 -norm regularization in the deep learning model controls the penalty for complexity, and we set the regularization parameter to $\alpha = 0.0001$. (2) early stoppage by using a validation set, resulting in a more conservative estimate of generalization. (3) the use of a smaller number of neurons in the hidden layers, thereby significantly reducing the number of parameters (weights) in the model. Each shape in the dataset is represented by a 400- or 900-dimensional geodesic feature vector, and we set the sizes of the hidden layers for the first and second autoencoders to $h_1 = 200$ and $h_2 = 100$, respectively. (4) the use of sparsity in the model helps downsize the active neurons. We set the target sparsity probability to $\rho = 0.15$ and the weight of sparsity regularization term to $\beta = 3$. In addition to dropout regularization which constrains the network adaptation to the data during training and hence improves its generalization ability, the use of cross-validation helps detect overfit models by determining how well the model generalizes to other datasets by partitioning the data. In all experiments, we partitioned our data into training, validation and testing sets in an effort to avoid the overfitting issues.

6. Conclusion

In this paper, we introduced a discriminative classifier using deep learning with geodesic moments. The proposed approach uses stacked sparse autoencoders with two hidden layers to learn high-level features, which were shown to offer a higher discrimination power for 3D shape classification. We showed through extensive experiments on several 3D shape benchmarks that our deep learning based approach substantially outperforms existing methods not only in terms of classification accuracy rates, but also in terms of standard error rates. For future work, we plan to apply the proposed approach to other 3D shape analysis problems. In addition to exploring pre-trained deep learning models on larger 3D shape benchmarks, we also plan to investigate other deep neural network architectures in an effort to further improve the classification accuracy results.

References

- [1] W. Shen, Y. Wang, X. Bai, H. Wang, L. Latecki, Shape clustering: common structure discovery, *Pattern Recognit.* 46 (2) (2013) 539–550.
- [2] S. Biasotti, A. Cerri, M. Abdelrahman, M. Aono, A. Ben Hamza, M. El-Melegy, A. Farag, V. Garro, A. Giachetti, D. Giorgi, A. Godil, C. Li, Y.-J. Liu, H. Martono, C. Sanada, A. Tatsuma, S. Velasco-Forero, C.-X. Xu, SHREC'14 track: retrieval and classification on textured 3d models, in: *Proceedings of the Eurographics Workshop on 3D Object Retrieval*, 2014, pp. 111–120.
- [3] C. Li, A. Stevens, C. Chen, Y. Pu, Z. Gan, L. Carin, Learning weight uncertainty with stochastic gradient MCMC for shape classification, in: *Proceedings of the CVPR*, 2016.
- [4] D. Pickup, X. Sun, P. Rosin, R. Martin, Z. Cheng, Z. Lian, M. Aono, A. Ben Hamza, A. Bronstein, M. Bronstein, S. Bu, U. Castellani, S. Cheng, V. Garro, A. Giachetti, A. Godil, L. Isaia, J. Han, H. Johan, L. Lai, B. Li, C. Li, H. Li, R. Litman, X. Liu, Z. Liu, Y. Lu, L. Sun, G. Tam, A. Tatsuma, J. Ye, Shape retrieval of non-rigid 3d human models, *Int J. Comput. Vision* 120 (2) (2016) 169–193.
- [5] M. Savva, F. Yu, H. Su, M. Aono, B. Chen, D. Cohen-Or, W. Deng, H. Su, S. Bai, X. Bai, J.H. N. Fish, E. Kalogerakis, E. Learned-Miller, Y. Li, M. Liao, S. Maji, Y. Wang, N. Zhang, Z. Zhou, SHREC'16 track: large-scale 3D shape retrieval from ShapeNet Core55, in: *Proceedings of the Eurographics Workshop on 3D Object Retrieval*, 2016.
- [6] R. Rostamov, Laplace-Beltrami eigenfunctions for deformation invariant shape representation, in: *Proceedings of the Symposium Geometry Processing*, 2007, pp. 225–233.
- [7] J. Sun, M. Ovsjanikov, L. Guibas, A concise and provably informative multi-scale signature based on heat diffusion, *Comput. Graphics Forum* 28 (5) (2009) 1383–1392.
- [8] M. Aubry, U. Schlickewei, D. Cremers, The wave kernel signature: a quantum mechanical approach to shape analysis, in: *Proceedings of the Computational Methods for the Innovative Design of Electrical Devices*, 2011, pp. 1626–1633.
- [9] M. Masoumi, A. Ben Hamza, A spectral graph wavelet approach for nonrigid 3D shape retrieval, *Pattern Recognit. Lett.* 83 (2016) 339–348.
- [10] M. Reuter, F. Wolter, N. Peinecke, Laplace-Beltrami spectra as 'Shape-DNA' of surfaces and solids, *Comput.-Aided Des.* 38 (4) (2006) 342–366.
- [11] A. Chaudhari, R. Leahy, B. Wise, N. Lane, R. Badawi, A. Joshi, Global point signature for shape analysis of carpal bones, *Phys. Med. Biol.* 59 (2014) 961–973.
- [12] Z. Gao, Z. Yu, X. Pang, A compact shape descriptor for triangular surface meshes, *Comput.-Aided Des.* 53 (2014) 62–69.
- [13] Z. Lian, A. Godil, B. Bustos, M. Daoudi, J. Hermans, S. Kawamura, Y. Kurita, G. Lavoué, H. Nguyen, R. Ohbuchi, Y. Ohkita, Y. Ohishi, F. Porikli, M. Reuter, I. Sipiran, D. Smeets, P. Suetens, H. Tabia, D. Vandermeulen, A comparison of methods for non-rigid 3D shape retrieval, *Pattern Recognit.* 46 (1) (2013) 449–461.
- [14] C. Li, A. Ben Hamza, Spatially aggregating spectral descriptors for nonrigid 3D shape retrieval: a comparative survey, *Multimedia Syst.* 20 (3) (2014) 253–281.
- [15] J. Schmidhuber, Deep learning in neural networks: an overview, *Neural Netw.* 61 (2015) 85–117.
- [16] Y. Bengio, Learning deep architectures for AI, *Found. Trends Mach. Learn.* 2 (1) (2009) 1–127.
- [17] S. Bu, Z. Liu, J. Han, J. Wu, R. Ji, Learning high-level feature by deep belief networks for 3-D model retrieval and recognition, *IEEE Trans. Multimedia* 24 (16) (2014) 2154–2167.
- [18] H. Su, S. Maji, E. Kalogerakis, E. Learned-Miller, Multi-view convolutional neural networks for 3D shape recognition, in: *Proceedings of the ICCV*, 2015, pp. 945–953.
- [19] C. Qi, H. Su, M. Nießner, A. Dai, M. Yan, L. Guibas, Volumetric and multi-view CNNs for object classification on 3D data, in: *Proceedings of the CVPR*, 2016.
- [20] Z. Zhu, X. Wang, S. Bai, C. Yao, X. Bai, Deep learning representation using autoencoder for 3D shape retrieval, *Neurocomputing* (2016).
- [21] S. Rosenberg, *The Laplacian on a Riemannian Manifold*, Cambridge University Press, 1997.
- [22] D. Aouada, H. Krim, Squigraphs for fine and compact modeling of 3D shapes, *IEEE Trans. Image Process.* 19 (2) (2010) 306–321.
- [23] O. Calin, D.-C. Chang, *Geometric Mechanics on Riemannian Manifolds: Applications to Partial Differential Equations*, Birkhäuser, 2005.
- [24] Z. Lian, A. Godil, B. Bustos, M. Daoudi, J. Hermans, S. Kawamura, Y. Kurita, G. Lavoué, H. Nguyen, R. Ohbuchi, Y. Ohkita, Y. Ohishi, F. Porikli, M. Reuter, I. Sipiran, D. Smeets, P. Suetens, H. Tabia, D. Vandermeulen, SHREC'11 track: shape retrieval on non-rigid 3D watertight meshes, in: *Proceedings of the Eurographics/ACM SIGGRAPH Symp. 3D Object Retrieval*, 2011, pp. 79–88.
- [25] Z. Lian, A. Godil, T. Fabry, T. Furuya, J. Hermans, R. Ohbuchi, C. Shu, D. Smeets, P. Suetens, D. Vandermeulen, S. Wuhrer, SHREC'10 track: non-rigid 3D shape retrieval, in: *Proceedings of the Eurographics/ACM SIGGRAPH Symp. 3D Object Retrieval*, 2010, pp. 101–108.
- [26] J.Z. Z. Lian, S. Choi, H. ElNaghy, J. El-Sana, T. Furuya, A. Giachetti, R.G.L. Isaia, L. Lai, C. Li, H. Li, F. Limberger, R. Martin, R. Nakanishi, A.N.L. Nonato, R. Ohbuchi, K. Pevzner, D. Pickup, P. Rosin, A. Sharf, L. Sun, X. Sun, S. Tari, G. Unal, R. Wilson, SHREC'15 track: non-rigid 3D shape retrieval, in: *Proceedings of the Eurographics Workshop on 3D Object Retrieval*, 2015, pp. 1–14.
- [27] M. Khabou, L. Hermi, M. Rhouma, Shape recognition using eigenvalues of the Dirichlet Laplacian, *Pattern Recognit.* 40 (2007) 141–153.
- [28] F. Hajati, A. Raie, Y. Gao, 2.5D face recognition using patch geodesic moments, *Pattern Recognit.* 45 (2012) 969–982.