Intra-class retrieval of non-rigid 3D objects: Application to Face Recognition

Georgios Passalis[†‡], Ioannis A. Kakadiaris[†], *Member, IEEE,* and Theoharis Theoharis[†‡]

***Abstract*— As the size of the available collections of 3D objects grows, database transactions become essential for their management, with the key operation being retrieval (query). Large collections are also pre-categorized into classes, so that a single class contains objects of the same type (e.g., human faces, cars, four-legged animals). It is shown that general object retrieval methods are inadequate for intra-class retrieval tasks. We advocate that such intra-class problems require a specialized method that can exploit the basic class characteristics in order to achieve higher accuracy. A novel 3D object retrieval method is presented which uses a Parameterized Annotated Model of the shape of the objects in a class incorporating its main characteristics. The annotated subdivision-based model is fitted onto objects of the class using a deformable model framework, converted to a geometry image and transformed into the wavelet domain. Object retrieval takes place in the wavelet domain. The method does not require user interaction, achieves high accuracy, it is efficient for use with large databases, and it is suitable for non-rigid object classes. We apply our method to the face recognition domain, one of the most challenging intra-class retrieval tasks. We utilized the Face Recognition Grand Challenge database, yielding an average verification rate of 95.2% at 10^{-3} false accept rate.**

***Index Terms*— H.3.3 Information Search and Retrieval, I.5.4d Face and gesture recognition**

I. Introduction

THREE dimensional (3D) object representations offer a new level of complexity and descriptive power compared to traditional 2D images. The introduction of cost effective 3D scanners has resulted in an explosion of the collection of available 3D objects. Database transactions on such collections, for example *categorization* (automatically dividing the database into meaningful categories for storage purposes) and *retrieval* (given a query object, search the database for similar objects) are thus becoming increasingly important. Retrieval is a key operation which must be performed accurately and efficiently for any large database to be useful. Moreover, it is usually the case that a large collection of objects will contain groups of similar objects, as a result of categorization. Such *intra-class*[1] retrieval problems more difficult, because of the similarity between the objects within a class.

Object retrieval can be applied to a large number of important tasks, making it one of the most widely researched topics. From web-based search engines, to security applications such as face recognition for access control, there is a need for efficient and robust object retrieval methods. To this end, many approaches have been proposed with the usual trade-off between retrieval accuracy and efficiency. These two goals often conflict as accuracy requires an object description that focuses on detail, while efficiency requires a more abstract and compact representation. In this paper, we focus on intra-class object retrieval problems, and more specifically on face recognition. By considering the human face as a class of objects, the task of verifying a person's identity can be expressed as an intra-class retrieval operation.

[†]CBL, Department of Computer Science, University of Houston.

[‡]CGL, Department of Informatics, University of Athens.

[1]In this paper, intra-class refers to retrieval on a database containing objects of the same class, while inter-class refers to retrieval on a database containing objects of multiple classes.

II. Related Work

A. Object Retrieval

Three dimensional object representation techniques can be divided into four categories according to a survey by Hlavaty *et al.* [1]: a) topology-based, b) histogram-based, c) view-based, and d) shape-based.

Topological methods include the work of Sundar *et al.* [2] who proposed the use of the skeleton of an object for object retrieval and the work of Hilaga *et al.* [3] who proposed the use of medial surfaces. Topological methods are the less suitable for intra-class object retrieval problems, since all instances of a single class have identical topology. The same holds true for histogram-based methods such as the ones proposed by Osada *et al.* [4] and Ankerst *et al.* [5]. Shape histograms are almost identical on classes with relatively homogeneous objects, such as human faces.

View-based methods [6], [7] extract 2D views of the objects and use them as features for retrieval purposes. There is an unavoidable loss of descriptive power when projecting 3D geometry onto a 2D view. In our previous work on object retrieval, we extended this idea to 3D views and presented a 3D rigid object retrieval method which employs depth buffers for representing and comparing the objects [8]. Specifically, multiple depth buffers per object (computed from different points of view) are compared for surface and volume similarity, inspired by our previous work on matching 3D complementary fragments of fractured objects [9].

Frome *et al.* [10] introduced two shape descriptors, 3D shape contexts and harmonic shape contexts, and used them in recognizing objects in a noisy range image. Kazhdan *et al.* [11], [12] used spherical harmonics as shape descriptors for 3D shape retrieval. Spherical harmonics is an attractive mathematical tool for this task due to its inherent rotationally invariant nature. Vranic [13] proposed several shape descriptors, including depth buffer-based, silhouette-based and ray-based. Each shape descriptor produces a feature vector, and

vectors from different objects are compared using L^1 or L^2 distances to determine similarity.

Laga *et al.* [14] proposed the use of geometry images for matching 3D objects. They applied the work of Hoppe *et al.* [15], [16] on geometry images to transform the 3D polygonal data to geometry images. They use the principal axis of the object to make their method rotationally invariant and evaluate it using a limited dataset (120 models). Although both this method and the proposed method utilize geometry images, they do not share the same goal, as the former is aimed at inter-class retrieval and the latter at intra-class retrieval.

The majority of the work conducted in 3D object retrieval concentrates on retrieving objects that belong to the same class while querying a multi-class database. Our work focuses on intra-class object retrieval, which has a different nature as a problem. The goal is not to describe the general shape that all objects within a class share so as to discern them from objects of a different class, but to find subtle differences and similarities among objects that share many common features. Therefore, even though both problems belong to the object retrieval domain, the methods that tackle them have different focus.

B. Face Recognition

Face recognition is one of the most challenging intra-class retrieval problems since human faces are non-rigid and relatively similar to each other. Also, face recognition has to compete with the accuracy of other biometric technologies. Hence the primary goal of face recognition is accuracy. Efficiency, which is the main goal in other retrieval applications, is secondary in face recognition.

Initially researchers concentrated on two-dimensional face recognition. Despite the introduction of commercial grade 2D face recognition systems, face recognition remains unreliable. Extensive experiments conducted using the FERET dataset [17], [18] and during the Face Recognition Vendor Test (FRVT) 2002 [19] study indicate that the success rate is not sufficient for critical applications. It appears that most 2D face recognition techniques stumble on the inherent problems of the 2D modality (e.g., pose, illumination).

Recently there has been an expansion of face recognition research to include 3D and multi-modal approaches due to the decrease of the cost of the acquisition hardware and expected gains in retrieval accuracy. Three-dimensional approaches utilize laser or optical (stereo vision) scanners to acquire the geometry of the human face, and multi-modal approaches combine different modalities (e.g., 2D, 3D and infrared). There is a constant research interest [20], [21], [22], [23] in all of these modalities. Excellent overviews of the research conducted on this field have been given by Chang *et al.* [24], [25], Bowyer *et al.* [26] and Zhao *et al.* [27].

In order to stimulate the development and evaluate the various face recognition approaches, the National Institute of Standards and Technology (NIST) set up the Face Recognition Grand Challenge (FRGC) in 2004 which aims to improve performance of face recognition systems by an order of magnitude [28]. Another FRVT [29] follows the conclusion of the Grand Challenge. In our previous work on face recognition, we made use of the extensive databases and well defined experiments provided by FRGC to evaluate our method [30], [31].

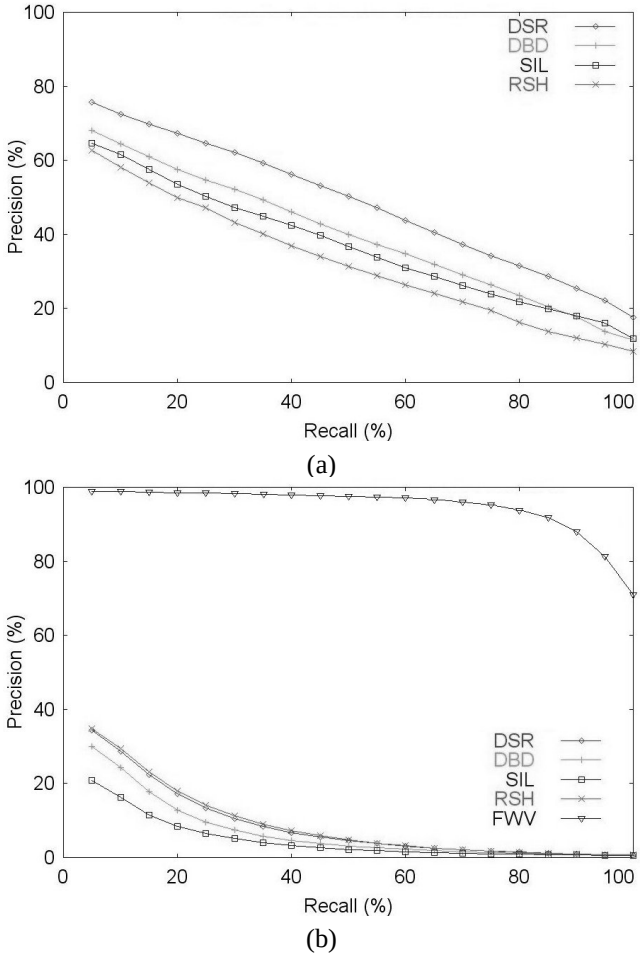

Fig. 1. Precision-Recall curves for (a) multi-class test database, (b) FRGC human face database. DSR, DBD, SIL and RSH refer to hybrid-based, depth buffer-based, silhouette-based and ray-based feature vectors respectively, provided by D. Vranic. FWV refers to our method.

III. Motivation

Our main assumption is that general object retrieval methods, even though they perform well in inter-class retrieval, are inadequate for intra-class retrieval, since they lack the necessary descriptive power. A method suitable for intra-class retrieval should take advantage of the characteristics of the class to achieve higher accuracy. The major difference compared to inter-class retrieval, is that an intra-class retrieval method needs to work only on a known class of objects, therefore it has the advantage of *a priori* knowledge about certain characteristics. The more of these characteristics that are exploited in the design and implementation of the method, the more capable the method becomes with respect to this class.

In order to validate our assumption we tested various general shape descriptors by using the tools available in Vranic's 3D search engine webpage [32], a detailed description of which can be found in [13]. We performed two retrieval experiments,

one with their test database, a multi-class database with 1841 miscellaneous objects, and one with the FRGC v2 human face database (a detailed description of which can be found in Sec. V-A). These methods, designed to operate on databases that contain objects from multiple classes, were selected as they are typical examples of general object retrieval methods and they are widely used as benchmarks in this domain. Moreover, their performance is considered challenging (Fig. 1(a)) when measured on typical databases that are used in the general object retrieval domain. But when their intra-class performance is measured on specialized and challenging databases their lack of descriptive power becomes evident (Fig.1 (b)). In Fig.1 (b) we also plot the performance of our method in the same intra-class experiment. These results are not meant for direct comparison; their purpose is to show that in an intra-class retrieval problem (such as face recognition), where general retrieval methods perform inadequately, it is possible to achieve high performance with a specialized intra-class method. In cases of very challenging retrieval problems it is necessary to sacrifice generality in order to achieve high performance.

In this paper, we generalize and extend our work on 3D face recognition [31], [33], by incorporating features from our previous 3D object retrieval work [8], thus making it suitable for intra-class retrieval in arbitrary 3D object classes. We then evaluate the accuracy and robustness of our method by utilizing the largest and most challenging publicly available 3D human face database.

The fundamental idea behind our method is to convert raw polygons in R^3 space into a compact 2D description that retains the geometry information, then perform the retrieval operation in R^2 space. This is advantageous because R^2 is simpler and several good 2D techniques exist. A 3D model is first created that describes the selected class. Apart from the geometry, the model also includes any additional features that are characteristic of the class (e.g., area annotation, landmarks). Additionally the model has a regularly sampled mapping from R^3 to R^2 (*UV parameterization*) that can be used to construct the equivalent 2D description, the geometry image [16], [34], [35]. Subsequently, a subdivision-based model is fitted onto all the objects of the class using a deformable model framework. The result is converted to a geometry image and a wavelet decomposition is applied. The wavelet coefficients are stored for matching and retrieval purposes.

The rest of the paper is organized as follows: Section IV describes our general method, Section V describes the application of our method to face recognition, Section VI discusses the advantages and limitations of our approach, while in Section VII we summarize our work.

IV. Methods

Our intra-class method combines the use of subdivision-based deformable models with geometry images. The idea is to reduce the dimensionality of the object representation while retaining the descriptive power of 3D polygonal data. The outline of our method follows:

1) Construct an annotated model. This step is performed once only for each class of objects.

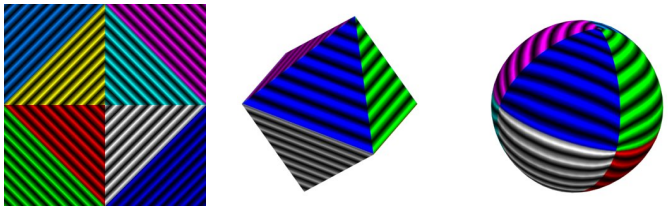

(a) (b) (c)

Fig. 2. Octahedron-based parameterization: (a) texture map in image space; (b) texture map projected onto octahedron and (c) texture map projected onto sphere.

2) **Preprocess:** For each object in the class:
   - Register object with the model.
   - Deform the model to fit the object using the Dynamic Subdivision Framework.
   - Convert the deformed model to a geometry image.
   - Apply wavelet analysis to the geometry image.
3) **Retrieval:** Perform object retrieval by comparing the wavelet coefficients.

Note that *model* refers to the annotated model that is deformed, and *object* refers to the polygonal object to which the model is fitted.

A. Annotated Models: Geometry Images and UV parameterization

In general, a UV parameterization denotes a mapping from R^3 to R^2. If a polygonal object has such a parameterization then each vertex is assigned a pair of (u, v) values. If this parameterization is injective it allows us to convert any polygonal object to an equivalent representation called geometry image [16], [34], [35]. A geometry image is a regularly sampled 2D image that has three channels, each one encoding geometric information (x,y and z components of a vertex in R^3). Because neighbor information is retained in the geometry image, each element (geometry pixel) can form a triangle in R^3 with two of its neighboring elements, thus allowing an easy transition back to the polygonal representation.

Praun and Hoppe [15] introduced spherical geometry images, which are suitable for genus[2] zero objects only. This topological limitation allows spherical geometry images to have several advantages over other geometry image types, such as continuity in image space, due to the lack of cuts. This approach takes advantage of the fact that an octahedron can be ”unfolded” to 2D space (Fig.2 (a,b)) and subsequently be subscribed to a sphere (Fig.2 (c)). Using this property, an injective mapping from a sphere in R^3 to a plane in R^2 is acquired. Finally, the object is mapped onto the sphere.

In our method, the UV parameterization needs to be computed only during the initial construction of the model. For topologically open models the problem of finding an efficient UV parameterization is relatively easy; texture mapping functions such as cylindrical mapping can be employed. For

[2]Genus is the topologically invariant property of a surface defined as the largest number of nonintersecting simple closed curves that can be drawn on the surface without separating it [36].

topologically closed, genus zero models, Praun's octahedron-based parameterization is utilized [15]. For closed models with higher genus, more advanced methods need to be applied such as the conformal parameterization proposed by Gu *et al.* [37].

Apart from the UV parameterization that is necessary for models of any class, there are optional attributes that can be assigned to the model. It is very useful to annotate the model into different areas that are associated with features of the specific class. This annotation is inherited by the objects through the process of fitting the model onto them, and can be used during the analysis of the results. Additionally, important landmarks can be assigned to model vertices, which will be positioned on the corresponding points on the object after fitting. The exact use of these attributes is application specific, but as shown in Sec. V, it can improve the object retrieval accuracy, especially when non-rigid objects are present.

Note that due to the flexible nature of subdivision surfaces (Sec. IV-C) there are few limitations to the topology or triangulation that the model must have. Almost any 3D polygonal mesh can be used as a basis for the model, and the additional attributes (e.g., UV parameterization and area annotation) can be subsequently added.

B. Registration

Before a polygonal object is converted to a geometry image through the fitting process, the object's orientation must be known. To this end, each object is registered with the annotated model, which has a known orientation. The Iterative Closest Point (ICP) algorithm [38] is the most commonly used algorithm for registration tasks. It determines correspondences between the vertices of the objects and minimizes the sum of the square of their distances. We employ an improvement suggested by Turk *et al.* [39] to reject vertex pairs containing points on surface boundaries. Additionally, in previous work we developed a registration algorithm [9] that uses a global optimization technique (Simulated Annealing - SA [40], [41]) applied to depth images. The objects are converted to depth images through OpenGL's z-buffers, and the discrete sum of differences over the z-buffer's derivatives is minimized.

Since each of these methods has its advantages and limitations, we utilize both in our method. The idea is that each algorithm will confute the shortcomings of the other. We currently use these algorithms sequentially, with ICP first because it is less sensitive to initial conditions, and SA second because it offers more fine registration. Obviously, the exact order in which these methods will be combined can vary, depending on the application.

Both of these algorithms perform better when applied on rigid objects. The reason is that object deformations cause different registration with respect to the non deformed version. To tackle this problem, we make use of the annotated model. During the registration we utilize only the areas that we expect to be rigid, disregarding any area that is subject to deformations. Our experiments showed that registration is one of the most important steps of our method and has a significant impact on accuracy.

Finally, note that we did not include any optimizations to lower the computational cost, except for the application of a space partitioning technique for ICP's nearest neighbor search. Currently, the registration process consumes one third of the total time of the preprocessing step. Therefore, our method's efficiency could benefit by optimized versions of the ICP algorithm such as the one proposed by Yan *et al.* [42].

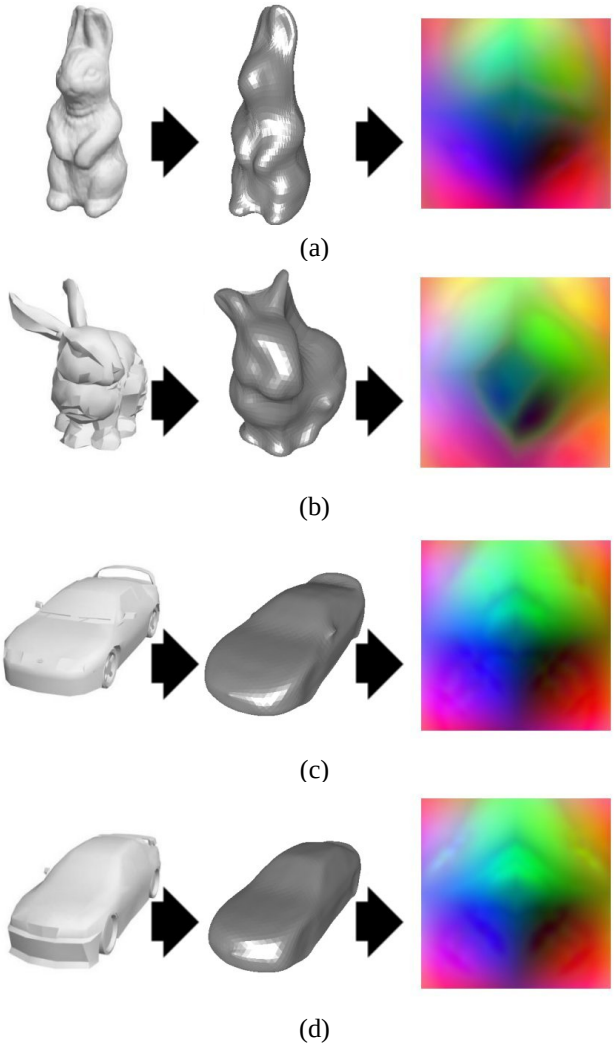

Fig. 3. Fitting example for two different genus zero classes (original object, fitted model and resulting geometry image): (a,b) rabbit class, (c,d) car class.

C. Dynamic Subdivision Framework

In our approach, we utilize the deformable model framework for data fitting purposes. There has been a constant research interest in deformable models since their introduction by Terzopoulos [43], [44]. The main idea behind this framework is that all the physical properties of a rigid or non-rigid object are analytically formulated and then an approximating solution is computed. These properties include mechanical properties (e.g., velocity, acceleration) and elastic properties of the surface of the object (e.g., strain energy, material stiffness). Recently Mandal *et al.* [45], [46] modified the deformable framework so that it uses subdivision surfaces instead of parametric surfaces. The analytical formulation remains the same but the Finite Element Method (FEM) implementation is different and it is adjusted for use with subdivision surfaces.

We selected the subdivision implementation because of the greater flexibility subdivision surfaces offer. Figure 3 depicts the application of our method on four polygonal objects, from two different classes. The model which was fitted in all cases was a sphere with an octahedron-based parameterization (Fig. 2 (c)).

1) Loop Subdivision Scheme: A subdivision surface [47] is a smooth explicit representation of the surface of an object. Compared to traditional parametric surfaces such as B-spline and Bezier, they offer more flexibility and scalability. A subdivision surface is defined by a polygonal mesh that is called control mesh and a set of rules that control how the subdivision behaves. Different rules define different schemes, with the most widespread being Doo-Sabin [48], Catmull-Clark [49], and Loop [50] schemes. A single subdivision step includes the splitting of each polygon to four smaller polygons and the repositioning of its vertices. The successive application of this subdivision step infinite times produces a smooth surface called limit surface.

In our implementation, we use the Loop scheme [50] which is simple, efficient, and offers significant advantages:

- it can be applied on triangular meshes with vertices of arbitrary valence;
- it produces a limit surface with C^2 smoothness; and
- the smooth patch produced for a single triangle depends only on the vertices inside the 1-neighborhood area of the vertices of the triangle.

2) Analytical Formulation: The deformation of a model is governed by the degrees of freedom assigned to it. Let $\mathbf{q}$ represent the n degrees of freedom of our model. In the case of subdivision surfaces this vector is the concatenation of the vertices of the control mesh. The equation that controls the deformations of the model is the following:

$$M\frac{d^2\mathbf{q}}{dt^2} + \mathbf{D}\frac{d\mathbf{q}}{dt} + \mathbf{K}\mathbf{q} = \mathbf{f_q}$$

where $\mathbf{M}$ is the $n \times n$ mass matrix, $\mathbf{D}$ the $n \times n$ damping matrix, $\mathbf{K}$ the $n \times n$ stiffness matrix and $\mathbf{f_q}$ is the $n \times 1$ vector of the external forces. The mass matrix is related to the kinetic energy of the model and the damping matrix to the energy dissipation.

In our implementation we used $\mathbf{M} = \emptyset$ and $\mathbf{D} = \emptyset$. The stiffness matrix is the most important component as it resists the external forces and determines elastic properties of the model. It can be decomposed into three matrices $\mathbf{K} = \mathbf{K}_{fo} + \mathbf{K}_{so} + \mathbf{K}_{sp}$. The matrix $\mathbf{K}_{fo}$ is related to the first order strain energy, $\mathbf{K}_{so}$ to the second order strain energy and $\mathbf{K}_{sp}$ is related to the spring forces energy:

$$E_{fo} = \frac{1}{2}\kappa_{fo}\mathbf{q}^T\mathbf{K}_{fo}\mathbf{q},$$

$$E_{so} = \frac{1}{2}\kappa_{so}\mathbf{q}^T\mathbf{K}_{so}\mathbf{q},$$

$$E_{sp} = \frac{1}{2}\kappa_{sp}\mathbf{q}^T\mathbf{K}_{sp}\mathbf{q}$$

where $\kappa_{fo}, \kappa_{so}, \kappa_{sp}$ are the individual weights.

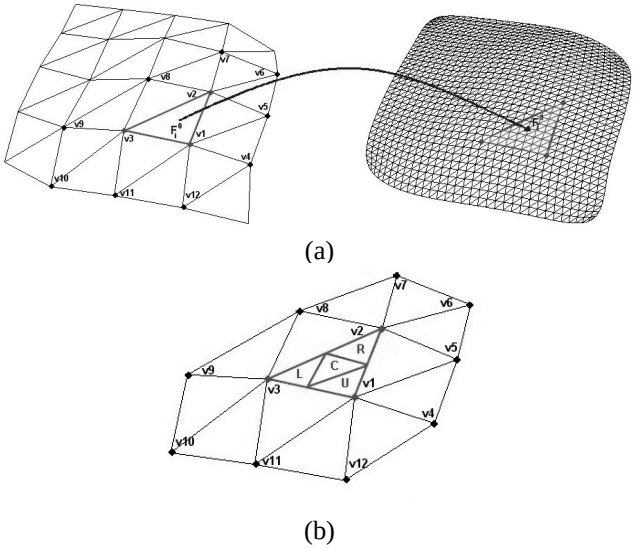

Fig. 4. (a) Subdivision surface: control points and limit surface. Finite element F_i^0 depends on vertices $\mathbf{v}_1 \ldots \mathbf{v}_{12}$ and its associated surface patch is F_i^3. (b) Up (U), left (L), right (R) and center (C) triangles produced by a single subdivision step.

3) Finite Element Method Implementation: The vertices in the control mesh of the subdivision surface determine the degrees of freedom of the model. The triangles in the control mesh are used as finite elements, with each triangle (denoted by F_i^0) controlling a triangular patch at the k^{th} level of subdivision (denoted by F_i^k) that is the approximation of the limit surface in that area. Every triangle inside the patch is used for computing the external forces and the stiffness matrix. Since the forces are actually applied in the control mesh and not in the limit surface, we need a way to map the forces from the limit surface back to the control mesh.

The answer to this problem lies in the way the triangular patch is computed by the Loop scheme. Suppose we have a mesh with valence six vertices only. Then, using the one-neighborhood property of the Loop scheme we need a total of 12 vertices to compute the shape of F_i^k independently of the actual value of k (Fig. 4(a)). We can create a vector $\mathbf{v}_i$ that is the concatenation of these 12 vertices. The shape function of F_i^k is given by:

$$\mathbf{s}_i^k(\mathbf{x}) = \mathbf{B}_i(\mathbf{x})\mathbf{v}_i$$

where $\mathbf{s}_i^k(\mathbf{x})$ is a 3×1 vector and $\mathbf{x} = (\alpha, \beta, \gamma)$ are the barycentric coordinates of the triangle in which $\mathbf{x}$ lies in the limit surface. The shape equation describes the points that lie in the limit surface with respect to the vertices in the control mesh. For a force computed in the limit surface we can use $\mathbf{B}_i(\mathbf{x})^T$ to map it to the control mesh.

The matrix $\mathbf{B}_i(\mathbf{x})$ is also easy to compute. In the first subdivision level, we have four 12×12 subdivision matrices $\mathbf{A}_U, \mathbf{A}_C, \mathbf{A}_L, \mathbf{A}_R$ that each transforms the vertices in v_i differently, depending on whether we want to produce the up, center, left or right triangle (Fig. 4(b)). For subsequent subdivision levels, we just multiply these matrices resulting in the matrix $\mathbf{A}_i^k$ from which we can construct the 36×36 matrix $\mathbf{A}c_i^k = \begin{pmatrix} \mathbf{A}_i^k & 0 & 0 \\ 0 & \mathbf{A}_i^k & 0 \\ 0 & 0 & \mathbf{A}_i^k \end{pmatrix}$. Finally, we multiply $\mathbf{A}c_i^k$ with

the 3×36 matrix $\mathbf{B}c(\mathbf{x})$:

$$\mathbf{B}c(\mathbf{x}) = \begin{pmatrix} \alpha & \beta & \gamma & .. & 0 & 0 & .. & .. & .. & 0 & 0 & .. & .. & .. & 0 \\ 0 & .. & .. & .. & 0 & \alpha & \beta & \gamma & .. & 0 & 0 & .. & .. & .. & 0 \\ 0 & .. & .. & .. & 0 & 0 & .. & .. & .. & 0 & \alpha & \beta & \gamma & .. & 0 \end{pmatrix}$$

to produce $\mathbf{B}_i(\mathbf{x})$. For each triangle in F_i^k we need to store the sequence that produced it (for example up, left, left, center) in order to compute $\mathbf{B}_i(\mathbf{x})$. Obviously $\mathbf{B}_i(\mathbf{x})$ changes depending on which triangle of F_i^k the point $\mathbf{x}$ lies.

With $\mathbf{B}_i(x)$ computed for each triangular patch, the computation of the stiffness matrix is straightforward. It is computed per each finite element and then the individual stiffness matrices are combined to create a global stiffness matrix. The energies corresponding to the three stiffness matrices for finite element F_i^k are the following:

$$E_{i,fo} = \frac{1}{2}\kappa_{sp}\mathbf{v}_i^T\mathbf{K}_{fo}\mathbf{v}_i,$$

$$E_{i,so} = \frac{1}{2}\kappa_{sp}\mathbf{v}_i^T\mathbf{K}_{so}\mathbf{v}_i,$$

$$E_{i,sp} = \frac{1}{2}\kappa_{sp}\mathbf{v}_i^T\mathbf{K}_{sp}\mathbf{v}_i$$

The matrices $\mathbf{K}_{sp}, \mathbf{K}_{fo}, \mathbf{K}_{so}$ can be derived from the $\mathbf{B}_i(x)$ as follows. Matrix $\mathbf{K}_{sp}$ is given by:

$$\sum_{\Omega}(\frac{|\mathbf{v}_a^k - \mathbf{v}_b^k| - l_{a,b}}{|\mathbf{v}_a^k - \mathbf{v}_b^k|})^2(\mathbf{B}_i(\mathbf{x}_a) - \mathbf{B}_i(\mathbf{x}_b))^T(\mathbf{B}_i(\mathbf{x}_a) - \mathbf{B}_i(\mathbf{x}_b))$$

where the domain Ω contains every pair of vertices $\mathbf{v}_a^k, \mathbf{v}_b^k$ that share an edge in F_i^k and $l_{a,b}$ is their distance before any deformation happens. Matrix $\mathbf{K}_{fo}$ is given by:

$$\sum_{\Omega}(\mathbf{B}_i(\mathbf{x}_a) - \mathbf{B}_i(\mathbf{x}_b))^T(\mathbf{B}_i(\mathbf{x}_a) - \mathbf{B}_i(\mathbf{x}_b))$$

where the domain Ω is the same as in $\mathbf{K}_{sp}$ and $\mathbf{K}_{so}$ is given by:

$$\sum_{\Omega'}(\mathbf{B}_i(\mathbf{x}_a) - 2\mathbf{B}_i(\mathbf{x}_b) + \mathbf{B}_i(\mathbf{x}_c))^T(\mathbf{B}_i(\mathbf{x}_a) - 2\mathbf{B}_i(\mathbf{x}_b) + \mathbf{B}_i(\mathbf{x}_c))$$

where domain Ω' contains all the triplets of vertices for all the triangles in F_i^k.

4) External Forces: External forces are the deformation driving forces. At each vertex in the limit surface the nearest neighbor on the object is found and this creates a deformation force, proportional to their distance. The nearest neighbor search is a costly operation, therefore a space partitioning technique (Octrees [51], [52]) is employed to improve performance.

In the FEM implementation the external forces are handled in a similar way with the stiffness matrix. A force applied on a vertex $\mathbf{v}_a^k$ of F_i^k when multiplied by $\mathbf{B}_i(\mathbf{x}_a)^T$ is distributed on the 12 vertices that control this element. These forces are then added to the global external forces vector. This property contributes to smoother results since even when a single force is applied to the limit surface, more than one control vertex is affected.

In order to improve the overall quality of fitting, we employ several filters on the nearest point selection that occurs during the computation of the external forces. Thus, we try to limit the external forces between points in the model and in the object that are likely to be erroneously corresponded.

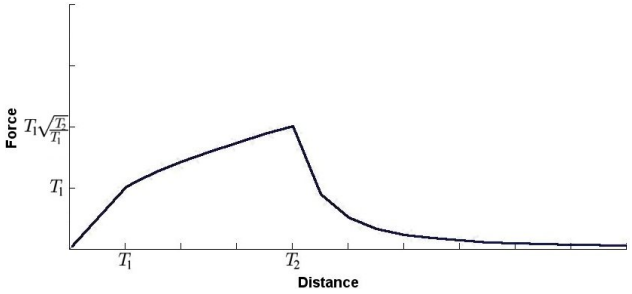

Fig. 5. Function employed for suppressing external forces.

We first use two filters based on the surface properties of the model and object, a normal and a curvature filter. The purpose of the filters is to disallow the matching of two points with different surface properties. For this purpose, the normal filter uses the dot product of the points' normals and compares it to a certain threshold. The curvature filter compares the absolute value of the difference of the points' curvature to a certain threshold. The exact curvature metric that will be used (e.g., K_1, K_2, Gaussian) is application specific. The final filter that is applied suppresses any external force between corresponding points with distance above a threshold. Normally, external forces are linearly proportional to the distance, but above a certain threshold we want the force to fade out, in order to prevent excessive deformations during a single iteration. The function used (visualized in Fig. 5), has two thresholds, T_1, T_2, that are application specific:

$$f(d) = \left\{ \begin{array}{cc} d, & if \quad d \leq T_1 \\ T_1\sqrt{d/T_1}, & if \quad d > T_1 \quad and \quad d \leq T_2 \\ \frac{T_1\sqrt{T_2/T_1}}{(d-T_2+1)^2}, & if \quad d > T_2 \end{array} \right\}$$

D. Wavelet Analysis

To compare our geometry images more robustly we performed wavelet analysis [53] using the Walsh transform. Each channel of the geometry image (X, Y and Z) is treated as a separate image for the wavelet analysis. The Walsh wavelet transform for images is a decimated wavelet decomposition using tensor products of the full Walsh wavelet packet system. The 1D Walsh wavelet packet system is constructed by repeated application of the Haar filter bank, a two-channel multirate filter bank based on the Haar conjugate mirror filter. Both channels output the result of convolving a 1D discrete signal with a Haar filter and then downsampling by a factor of two. The low–pass and high–pass Haar filters are g and h, respectively: $g = \frac{1}{\sqrt{2}}[1 \quad 1]$ and $h = \frac{1}{\sqrt{2}}[1 \quad -1]$.

For images, we use tensor products of these 1D filters. This means that the filter bank operations are applied separately to the rows and columns of the image, resulting in a four channel filter bank with channels LL, LH, HL, and HH (corresponding to the filters $g^t * g, g^t * h, h^t * g$ and $h^t * h$ respectively). In other words, channel LL (low–pass) captures the local averages of the image, and channels LH, HL and HH (high–pass) capture horizontal, vertical and diagonal edges, respectively. We recursively apply this decomposition to each of the four output channels to construct the full Walsh wavelet packet

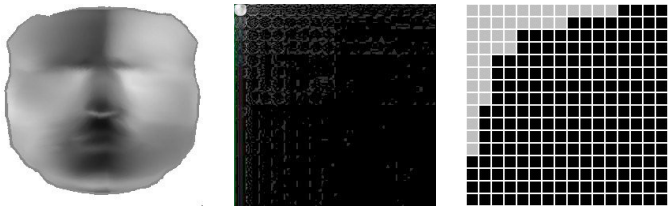

(a) (b) (c)

Fig. 6. Wavelet analysis of a facial geometry image: (a) Original image; (b) Four level Walsh transform; (c) Mask that selects 15% of wavelet bands.

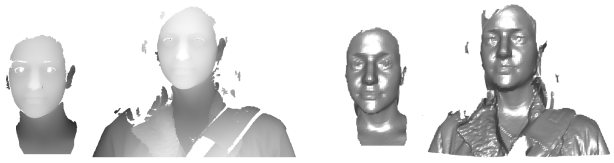

(a) (b)

Fig. 7. Data from FRGC v2. (a) Range data and (b) computed 3D meshes from a single individual acquired at different times.

tree decomposition. Conjugate mirror filter banks achieve perfect reconstruction, so the Walsh transform preserves all information originally present in the signal.

Due to the fact that we utilize area weighting (see Sec.IV-E), we selected the Haar filter as we value more its excellent localization properties rather than the better spectral characteristics of more complex wavelet transforms. The same reason dictated the use of the full packet decomposition over the traditional transform (where it is recursively applied only on the low–pass band at each level). The packet decomposition produces bands of equal resolution, allowing a straightforward application of the area weights. Additionally, each band has the same number of wavelet coefficients making the normalization of their weights unnecessary.

In order to make retrieval more efficient we keep only a subset of the coefficients. We utilize a mask that selects specific bands from the decomposition. The mask is the same for all objects allowing us to compare the coefficients without the need to reconstruct the geometry image. The criteria we used in order to construct the mask were the following:

1) The mask should retain most of the L^1 energy of the geometry image (thus favoring the low–pass bands).
2) The mask should exclude bands that are subject to noise.
3) The mask must be symmetrical with respect to the vertical and horizontal directions.

An example of this wavelet analysis (original image, decomposition, selection mask) from the Face Recognition domain is depicted in Fig. 6. Note that selecting coefficients from a wavelet packet decomposition has been used in the past for image retrieval purposes (more specifically, in the field of 2D Face Recognition [54]). In that case however, the nature of the data is different and the wavelet decomposition needs to address different problems (e.g., 2D images suffer from illumination problems while geometry images do not). This results in completely different coefficient selection.

E. Distance Metric

In order to perform object retrieval we need a way to compare the final data produced by our method (i.e., the wavelet coefficients). Since the wavelet coefficients are organized as a 2D image, object retrieval becomes an image retrieval problem, allowing us to use an L^1 norm in image space. In addition, by using the annotation of the model, we segment the geometry image into different areas and assign to each area F_k a different weight w_k. The weights can either be set statically from the *a priori* knowledge we have about this class, or they can be computed dynamically from characteristics of the specific object. Thus, in the X component (first geometry channel) we compute the L^1 norm of the difference:

$$\text{Score}_x(F_k) = \int_{F_k} |\text{Probe}_x(u,v) - \text{Gallery}_x(u,v)| du dv$$

and similarly for Score_y and Score_z. This score can be computed as a discrete sum over the stored wavelet coefficients because of the very good localization properties of Haar wavelets. Thus, for each area F_k we compute:

$$\text{Score}_x(F_k) = \sum_{i,j} |\text{Probe}_x[i,j] - \text{Gallery}_x[i,j]|$$

where the sum is over all coefficients $[i,j]$ that correspond to area F_k, and similarly for Score_y and Score_z. The total score for F_k is given by the sum:

$$\text{Score}(F_k) = \text{Score}_x(F_k) + \text{Score}_y(F_k) + \text{Score}_z(F_k)$$

while the score for the whole geometry image is:

$$\text{Score} = \sum_k w_k \text{Score}(F_k)$$

.

V. Application to Face Recognition

Face recognition, as one of the more challenging object retrieval problems, is suitable for evaluating the accuracy of novel retrieval methods. It is an intra-class retrieval problem, with non-rigid objects because of facial expressions. A face recognition system must be able to match a given instance (3D facial object) of a specific person with another instance of the same person stored in a database along with instances from many other people. Our general method was applied to the face recognition problem using the databases and experiments provided by the Face Recognition Grand Challenge (FRGC) [28].

A. FRGC v2 database

The FRGC provides data that consist of 3D images (range data) and high resolution controlled and uncontrolled stills. In this paper, we use the data from FRGC database v2, a total of 4007 range images (e.g., Fig. 7(a)), acquired between 2003 and 2004. The hardware used to acquire these range data was a Minolta Vivid 900 range scanner, with a resolution of 640x480. These data were obtained from 466 subjects and contain various facial expressions (e.g., happiness, surprise). The subjects are 57% male and 43% female, while the age

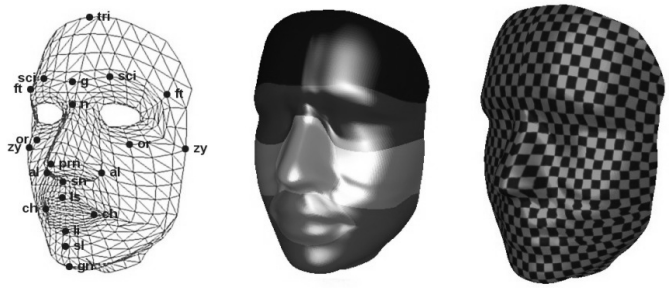

(a) (b) (c)

Fig. 8. (a) Anthropometric landmarks used; (b) Segmentation into annotated areas; and (c) Checkerboard texture to demonstrate parameterization.

distribution is 65% 18-22 years old, 18% 23-27 and 17% 28 or over.

The range data of the FRGC v2 database are converted to polygonal form (Fig. 7(b)). Before this conversion takes place, we apply a median cut filter for de-noising purposes. Note that the resolution of the geometry images is completely independent of the resolution of the input data, as the geometry images are obtained by regular sampling of the deformed model in UV space.

B. Annotated Face Model

A human face model was constructed according to the principles described in Sec. IV-C. The Annotated Face Model (AFM) is based on an average facial 3D mesh, constructed using statistical data. Anthropometric landmarks are associated with its vertices (Fig. 8 (a)) based on the seminal work of Farkas [55]. Using information from facial physiology, we have annotated the AFM into different areas. Fig. 8(b) depicts our annotation with each area denoted by a different shade. These facial areas have different properties associated with them, for example the mouth area is expected to be non-rigid, while the nose area is rigid in all scans of the same human face. The relatively simple topology of the human face allowed the use of a spherical mapping function to apply a continuous global UV parameterization. In addition to the necessary injective property, this mapping has the additional property that the ratio of area in 3D space to area in UV space is approximately the same for every part of the surface, leading to low distortion (Fig. 8(c)).

C. Results

We have planned our results to conform to standard tests used in face recognition, so that the performance of our method is directly comparable.

1) Identification: In the identification scenario, we have a 3D dataset for each person in the gallery set, while in the probe set there may be more than one instance of each person. This guarantees that each subject in the probe set will have exactly one match in the gallery set, resulting in 466 images in the gallery set and 3541 in the probe set (gallery/probe ratio is $\frac{2}{15}$). Note that this ad hoc division does not guarantee that the data in the gallery set will have a neutral expression, resulting in the existence of facial expressions in both gallery and probe sets, making the experiment challenging.

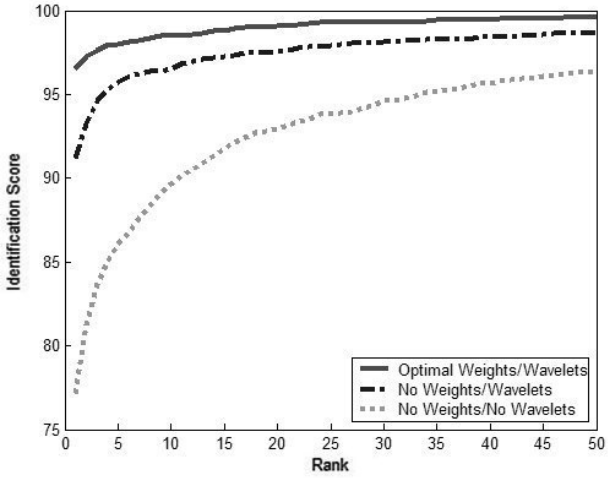

Fig. 9. CMC curves for identification experiment using the FRGC v2 database.

TABLE I
VERIFICATION RATES FOR ROC I, II, AND III FRGC EXPERIMENTS.

	FAR	Our method	Baseline
ROC I	10^{-3}	95.8%	56.2%
ROC I	10^{-2}	97.9%	78.4%
ROC II	10^{-3}	95.3%	49.5%
ROC II	10^{-2}	97.6%	75.7%
ROC III	10^{-3}	94.7%	42.7%
ROC III	10^{-2}	97.2%	72.5%

We evaluate the accuracy of our method using a Cumulative Match Characteristic (CMC) curve. Making use of all our method's features we achieved a rank one rate of 96.5%. In order to evaluate the importance of certain features (the area weights and the wavelet transform), we also report the performance without them (Fig. 9). If we disable the area weights the rank one rate drops to 91.1%, while if we additionally disable the wavelet transform and directly use the geometry image the rank one rate drops to 77.1%. The wavelet transform allows us to use only these coefficients which describe useful geometric features.

2) Verification: The FRGC, apart from the database, provides a set of experiments along with a baseline algorithm to allow researchers to evaluate their results. These experiments fall under the verification scenario, where the verification rate is measured for given false accept rates (FAR) and the results are summarized using Receiver Operating Characteristic (ROC) curves. Although a verification scenario is not the usual way to evaluate a general purpose object retrieval system, we present our results in this format so that direct comparisons with other face recognition approaches can be made.

In the FRGC, three masks are defined over the square similarity matrix which holds the computed similarity values between every pair of facial records. Each mask is used to perform a different verification experiment, thus producing

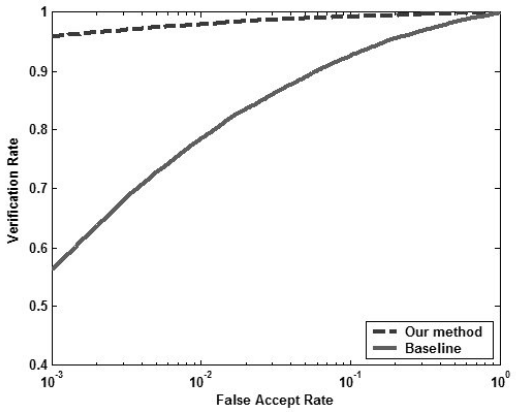

(a)

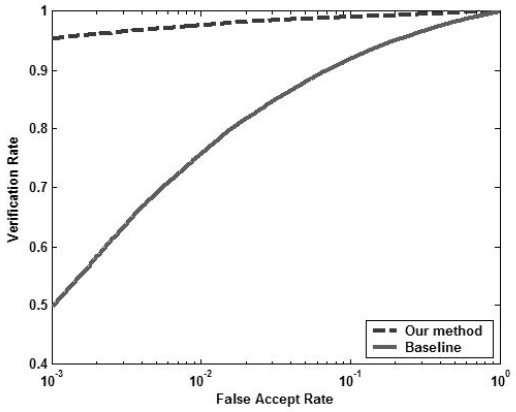

(b)

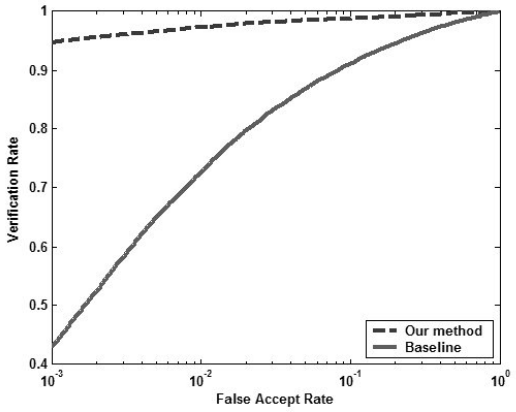

(c)

Fig. 10. Verification scenario: Our method versus the FRGC2 3D-only baseline. (a) ROC I, (b) ROC II, (c) ROC III.

three different ROC curves, which will be referred to as ROC I, II, and III. A mask selects a subset of the records to be used as the gallery set and another subset to be used as the probe set. In ROC I the utilized subset contains facial data that were captured within a semester, in ROC II it contains data that were captured within a year, while in ROC III the data are between semesters. These experiments are of increasing difficulty.

In Fig. 10 we compare the performance of our algorithm with the FRGC v2 3D-only baseline using the three default ROC curves, on the full database. By comparing the verification rates (Table I) we can deduce that our method outperforms

Fig. 11. Examples of challenging data from the FRGC v2 database.

the 3D-only baseline algorithm by 46% at 10^{-3} and by 22% at 10^{-2} on average. The FRGC baseline algorithms are based on principal component analysis (PCA) [56] optimized for large scale problems, and the distance in the nearest neighbor classifier is Whitened cosine. The matching algorithms for 3D data are based on the work of Chang *et al.* [57]. The FRGC baseline algorithm is included not as the most competitive approach, but as its name implies, as a baseline. Compared to other approaches in the field, the performance of our method is the highest reported for the FRGC v2 database (according to the latest survey of Chang *et al.* [25]). Finally, note that for these results we used an L^1 metric; using an L^2 metric results on an average performance drop of 0.1% at 10^{-3}.

VI. Discussion

From an accuracy point of view, we show that intra-class retrieval problems such as face recognition require specialized methods and not general object retrieval methods. Our results in face recognition demonstrate the accuracy of our method in a very challenging database (Fig. 11). Moreover, it becomes apparent that an inherent advantage of our method is that we can utilize the annotation of the model to increase performance. This is shown in the identification scenario, where by disregarding the annotation (and the corresponding weights) there was a significant decrease in performance (Fig. 9). Similarly, in the verification scenario, one of the reasons that we significantly outperform the baseline algorithm (Fig. 10) is that the baseline algorithm uses information from the entire face and does not disregard unreliable and non-rigid facial areas.

A detailed error analysis in the face recognition experiments revealed that the majority of the failure cases are related to the registration step. This happens because the output of the registration step unavoidably affects the fitting of the model, and the final representation (geometry images) is not rotationally invariant. In cases of intense facial expressions, the registration quality is worse than neutral expressions, even if both registration methods (ICP and SA) find the global minimum for their optimization functions.

From an efficiency point of view, the proposed method is suitable for large databases. It has a relatively heavy preprocessing step for each new object entering the database but has a very efficient retrieval step. In the face recognition experiments, the average preprocessing time for a facial scan (containing several thousand polygons) was approximately 20sec. The average retrieval rate was 10000 subjects per second. These measurements were carried out on a typical PC (3Ghz CPU/1GB RAM). The retrieval step is efficient because the preprocessing step produces a compact representation for

each object that stores the wavelet coefficients sequentially in a binary file. The comparison is then carried out on these coefficients directly, without the need to reverse the wavelet transform.

The efficiency of our method can be significantly improved by implementing a progressive approach. The wavelet packet decomposition has a mutli-resolution nature which can be used for this purpose. In our current approach, we utilize coefficients only from the last level of the decomposition. Instead, we can keep coefficients from other levels (especially from the low-pass bands) that will describe only the object's main geometric characteristics. The resulting reduced set of coefficients can be used for indexing purposes, and only objects similar enough will be compared using the full coefficient set.

Finally, the proposed method can be applied to classes of objects with varying resolutions. Neither its efficiency nor its accuracy are significantly affected, since the fitting of the model effectively re-samples the object. This is particularly important in practical situations, such as face recognition applications, where the data may be captured by different equipment at various resolutions.

VII. Conclusion

A method suitable for intra-class 3D object retrieval that focuses on the detailed representation of objects with similar shape and topology was presented. Using a model-based approach, 3D polygonal objects are converted to geometry images allowing one to take advantage of the efficiency of 2D representations while retaining the descriptive power of 3D data. Moreover it is shown that the careful annotation of rigid and non-rigid areas of the model can contribute to increased performance. The proposed method is evaluated through its application in 3D face recognition. By using the FRGC v2 database, the largest and most established 3D facial dataset currently available, useful conclusions about its performance and limitations are drawn. Finally, the compact wavelet representation ensures that the computational cost is minimized during object retrieval making the proposed approach suitable for real life applications where very large databases are used.

References

[1] T. Hlavaty and V. Skala, "A survey of methods for 3D model feature extraction," Seminar on Geometry and Graphics in Teaching Contemporary Engineer.

[2] H. Sundar, D. Silver, N. Gagvani, and S. Dickinson, "Skeleton based shape matching and retrieval," in *Proc. of Shape Modelling and Application Conference*, Seoul, S. Korea, May 12-15 2003, pp. 130–142.

[3] M. Hilaga, Y. Shinagawa, T. Kohmura, and T. Kunii, "Topology matching for fully automatic similarity estimation of 3D shapes," in *Proc. of SIGGRAPH*, Los Angeles, CA, 2001, pp. 203–212.

[4] R. Osada, T. Funkhauser, B. Chazalle, and D. Dobkins, "Matching 3D models with shape distribution," *ACM Transactions on Graphics*, vol. 21, no. 4, pp. 807–832, October 2002.

[5] M. Ankerst, G. Kastenm̈uller, H. P. Kriegel, , and T. Seidl, "3D shape histograms for similarity search and classification in spatial databases," in *Proc. of 6th International Symposium on Spatial Databases*, Hong-Kong, China, July 1999, pp. 207–226.

[6] T. Krüger, J. Wickel, P. Alvarado, and K. Kraiss, "Feature extraction from VRML models for view-based object recognition," in *Proc. of the 4th European Workshop on Image Analysis for Multimedia Interactive Services, 2003*, London, April 9-11 2003, pp. 391–394.

[7] D. Chen, X. Tian, Y. Shen, and M. Ouhyoung, "On visual similarity based 3D model retrieval," *EUROGRAPHICS 2003*, vol. 22, no. 3, pp. 223–232, 2003.

[8] N. Vajramushti, I. A. Kakadiaris, T. Theoharis, and G. Papaioannou, "Efficient 3D object retrieval using depth images," in *Proc. of the 6th ACM SIGMM International Workshop on Multimedia information*, New York, USA, October 2004, pp. 189–196.

[9] G. Papaioannou, E. Karabassi, and T. Theoharis, "Reconstruction of three-dimensional objects through matching of their parts," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 24, no. 1, pp. 114–124, January 2002.

[10] A. Frome, D. Huber, R. Kolluri, T. Bulow, and J. Malik, "Recognizing objects in range data using regional point descriptors," in *Proc. of the European Conference on Computer Vision*, May 2004, pp. 224–237.

[11] M. Kazhdan, "Shape representations and algorithms for 3d model retrieval," Ph.D. dissertation, Princeton University, June 2004.

[12] M. Kazhdan, T. Funkhouser, and S. Rusinkiewicz, "Rotation invariant spherical harmonic representation of 3D shape descriptors," in *Proc. of Eurographics Symposium on Geometry Processing*, Aachen, Germany, June 2003.

[13] D. Vranic, "3D model retrieval," Ph.D. dissertation, Universitat Leipzig, May 2004.

[14] H. Laga, H. Takahashi, and M. Nakajima, "Geometry image matching for similarity estimation of 3D shapes," in *IEEE Proc. of Computer Graphics International*, Crete, Greece, June 2004, pp. 490–496.

[15] E. Praun and H. Hoppe, "Spherical parametrization and remeshing," in *Proc. of SIGGRAPH*, San Diego, CA, July 2003, pp. 340–349.

[16] X. Gu, S. Gortler, and H. Hoppe, "Geometry images," in *Proc. of SIGGRAPH*, San Antonio, TX, July 2002, pp. 355–361.

[17] J. Phillips, H. Moon, S. Rizvi, and P. Rauss, "The FERET evaluation methodology for face-recognition algorithms," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 22, no. 10, pp. 1090–1104, October 2000.

[18] P. Phillips, A. Martin, C. Wilson, and M. Przybocki, "An introduction to evaluating biometric systems," *IEEE Computer*, vol. 33, no. 2, pp. 56–63, Feb. 2000.

[19] J. Phillips, P. Grother, R. Michaels, D. Blackburn, E. Tabassi, and J. Bone, "FRVT 2002: overview and summary," March 2003, www.frvt.org.

[20] K. Chang, K. Bowyer, and P. Flynn, "ARMS: Adaptive rigid milti-region selection for handling expression variation in 3D face recognition," in *IEEE Proc. of Workshop on Face Recognition Grand Challenge*, San Diego, CA, June 2005.

[21] V. Blanz, P. Grother, J. Phillips, and T. Vetter, "Face recognition based on frontal views generated from non-frontal images," in *IEEE Proc. of Computer Vision and Pattern Recognition*, San Diego, CA, June 2005.

[22] X. Liu and T. Chen, "Pose-robust face recognition using geometry assisted probabilistic modeling," in *IEEE Proc. of Computer Vision and Pattern Recognition*, San Diego, CA, June 2005.

[23] J. Heo, M. Savvides, and B. Vijayakumar, "Performance evaluation of face recognition using visual and thermal imagery with advanced correlation filters," in *IEEE Proc. of Computer Vision and Pattern Recognition*, San Diego, CA, June 2005, p. 9.

[24] K. I. Chang, K. W. Bowyer, and P. J. Flynn, "An evaluation of multimodal 2D+3D face biometrics," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 27(4), pp. 619–624, April 2005.

[25] K. Chang, K. Bowyer, and P. Flynn, "A survey of approaches and challenges in 3D and multi-modal 2D+3D face recognition," *Computer Vision and Image Understanding*, vol. 101, no. 1, pp. 1–15, January 2006.

[26] K. Bowyer, K. Chang, and P. Flynn, "A survey of approaches to 3D and multi-modal 3D+2D face recognition," in *IEEE Proc. of Int'l Conf. on Pattern Recognition*, Cambridge, UK, Aug. 2004, pp. 358–361.

[27] W. Zhao, R. Chellappa, P. J. Phillips, and A. Rosenfeld, "Face recognition: A literature survey," *ACM Comput. Surv.*, vol. 35, no. 4, pp. 399–458, 2003.

[28] P. Phillips, P. Flynn, T. Scruggs, K. W. Bowyer, J. Chang, K. Hoffman, J. Marques, J. Min, and W. Worek, "Overview of the Face Recognition Grand Challenge," in *IEEE Proc. of Computer Vision and Pattern Recognition*, San Diego, CA, June 2005.

[29] "Face Recognition Vendor Test 2006," http://www.frvt.org/FRVT2006/.

[30] I. Kakadiaris, G. Passalis, T. Theoharis, G. Toderici, I. Konstantinidis, and N. Murtuza, "Multimodal face recognition: Combination of geometry with physiological information," in *IEEE Proc. of Computer Vision and Pattern Recognition*, San Diego, CA, June 20 - 26 2005, pp. 1022–1029.

[31] G. Passalis, I. Kakadiaris, T. Theoharis, and N. Murtuza, "Evaluation of UR3D: Using an annotated deformable model for 3D face recognition," in *IEEE Proc. of Workshop on Face Recognition Grand Challenge*, San Diego, CA, June 2005.

[32] D. Vranic, "Content-based classification of 3D-models by capturing spatial characteristics," http://merkur01.inf.uni-konstanz.de/CCCC/.

[33] G. Passalis, "Three-dimensional face recognition," Master's thesis, University of Houston, May 2004.

[34] I. Kakadiaris, L. Shen, M. Papadakis, D. Kouri, and D. Hoffman, "g-HDAF multiresolution deformable models for shape modeling and reconstruction," in *British Machine Vision Conference*, Cardiff, UK, September 2-5 2002.

[35] I. Kakadiaris, M. Papadakis, L. Shen, D. Kouri, and D. Hoffman, "m-HDAF multiresolution deformable models," in *Proc. of the 14th International Conference on Digital Signal Processing*, Santorini, Greece, July 1-3 2002, pp. 505–508.

[36] "Mathworld," http://mathworld.wolfram.com/.

[37] X. Gu and S. Yau, "Global conformal surface parameterization," in *Proc. of Eurographics Symposium on Geometry Processing*, June 2003, pp. 127–137.

[38] P. Besl and N. McKay, "A method for registration of 3-D shapes," *IEEE Trans. on Pattern Analysis and Machine Intellignence*, vol. 14, no. 2, pp. 239–256, Feb. 1992.

[39] G. Turk and M. Levoy, "Zippered polygon meshes from range images," in *Proc. of SIGGRAPH*, Orlando, FL, July 1994, pp. 311–318.

[40] S. Kirkpatrick, C. Gelatt, and M. Vecchi, "Optimization by simulated annealing," *Science*, vol. 22, no. 4598, pp. 671–680, 1983.

[41] P. Siarry, G. Berthiau, F. Durbin, and J. Haussy, "Enhanced simulated annealing for globally minimizing functions of many-continuous variables," *ACM Transactions on Mathematical Software*, vol. 23, no. 2, pp. 209–228, 1997.

[42] P. Yan and K. Bowyer, "A fast algorithm for ICP-based 3D shape biometrics," in *Fourth IEEE Workshop on Automatic Identification Advanced Technologies (AutoID)*, October 2005.

[43] D. Terzopoulos, J. Platt, A. Barr, and K. Fleischer, "Elastically deformable models," in *Proc. of SIGGRAPH*, vol. 21(4), 1987, pp. 205–214.

[44] D. Terzopoulos and K. Fleischer, "Deformable models," *The Visual Computer*, vol. 4(6), pp. 306–331, 1988.

[45] C. Mandal, H. Qin, and B. Vemuri, "Dynamic smooth subdivision surfaces for data visualization," in *IEEE Proc. of Visualization*, Phoenix, Arizona, October 1997, pp. 371–377.

[46] C. Mandal, "A dynamic framework for subdivision surfaces," Ph.D. dissertation, University of Florida, 1998.

[47] D. Zorin and P. Schroeder, "Subdivision for modeling and animation," in *SIGGRAPH Course Notes*, 2000.

[48] D. Doo and M. Sabin, "Analysis of the behaviour of recursive division surfaces near extraordinary points," *Computer Aided Design*, vol. 10(6), pp. 356–360, 1978.

[49] E. Catmull and J. Clark, "Recursive generated B-spline surfaces on arbitrary topological meshes," *Computer Aided Design*, vol. 10(6), pp. 350–355, 1978.

[50] C. Loop, "Smooth subdivision surfaces based on triangles," Master's thesis, Department of Mathematics, University of Utah, 1987.

[51] A. Klinger, "Pattern and search statistics," *Optimizing Methods in Statistics*, pp. 303–337, 1971.

[52] A. Finkel and J. Bentley, "Quad trees. a data structure for retrieval of composite keys," *Acta Informatica*, vol. 4, no. 1, pp. 1–9, 1974.

[53] E. Stollnitz, T. DeRose, and D. Salesin, *Wavelets for Computer Graphics: Theory and Applications*. Morgan Kaufmann Publishers, Inc, 1996.

[54] R. Bhagavatula and M.Savvides, "PCA vs. automatically pruned wavelet-packet PCA for illumination tolerant face recognition," in *Fourth IEEE Workshop on Automatic Identification Advanced Technologies (AutoID)*, October 2005.

[55] L. Farkas, *Anthropometry of the Head and Face*. Raven Press, 1994.

[56] M. Turk and A. Pentland, "Eigenfaces for recognition," *J. of Cognitive Neuroscience*, vol. 3, no. 1, pp. 71–86, 1991.

[57] K. Chang, K. Bowyer, and P. Flynn, "Face recognition using 2D and 3D facial data," in *Proc. of Workshop on Multimodal User Authentication*, Dec. 2003, pp. 25–32.

Georgios Passalis, PhD Candidate.

Georgios has received his Bachelor from the Department of Informatics and Telecommunications of University of Athens. Subsequently he received his Master's degree from the Department of Computer Science of University of Houston. Currently his is a PhD Candidate in University of Athens. His thesis is focused on the scientific domains of Computer Graphics and Computer Vision. His research interests include object retrieval, face recognition, hardware accelerated voxelization and object reconstruction.

Ioannis A. Kakadiaris, Associate Professor.

Ioannis is the founder and director of UH's Computational Biomedicine Lab (formerly the Visual Computing Lab) and the Director of the Division of Bioimaging and Biocomputation at the UH Institute for Digital Informatics and Analysis. He is the recipient of the year 2000 NSF Early Career Development Award, UH Computer Science Research Excellence Award, UH Enron Teaching Excellence Award, James Muller VP Young Investigator Prize, and the Schlumberger Technical Foundation Award.

Theoharis Theoharis, Associate Professor.

Theoharis, received his D.Phil. in computer graphics and parallel processing from the University of Oxford, U.K. in 1988. He subsequently served as a research fellow (postdoc) at the University of Cambridge, U.K. and as a consultant with Andersen Consulting. He is with the University of Athens, Greece since 1993. During the 2002/3 academic year he was on Sabbatical leave at the Department of Computer Science, University of Houston, Texas. His main research interests lie in the fields of Computer Graphics, Visualization, Biometrics and Archaeological Reconstruction.