PHOG:Photometric and geometric functions for textured shape retrieval

Presentation by Eivind Kvissel

Introduction

This paper is about tackling the issue of textured 3D object retrieval.

Thanks to advances in modeling techniques, and to cheaper, yet effective 3D acquisition devices, there has been a dramatic increase in the amount of 3D sets available. This has lead to developments of 3D search engines able to organize and retrieve the 3D content, by using 3D descriptors and similar measures.

However, most descriptors proposed thus far are confined to shape, that is, they only analyze the geometric and/or topological properties of the given 3D objects.

Introduction

These days, many sensors are able to acquire not only the 3D shape, but also the texture of the object, and modern multiple-view stereo techniques enable the recovery of both geometric and photometric information from images.

And so we come to the purpose of this research paper. It proposes a shape description and retrieval pipeline which is able to handle both geometric and photometric information

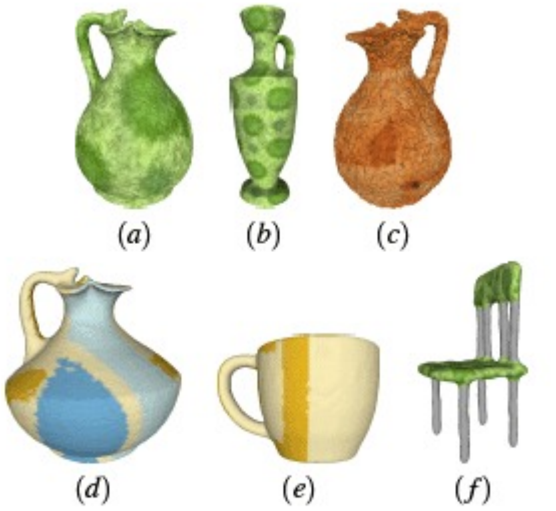

Some examples from the SHREC' 13 dataset used for testing

Methods Introduction

This paper proposes a shape description and retrieval pipeline which is able to handle both geometric and photometric information.

The descriptor described within consists of three parts:

- A purely photometric descriptor
  - Uses coordinates in the CIELab colour space to get persistence diagrams and spaces, which again is used to compute an approximation of persistence spaces
- An hybrid descriptor
  - Jointly captures shape and textures by persistent diagrams working on a generalization of the geodesic distance which fuses shape and color properties
- A purely geometric descriptor
  - Adopts a clustering technique on the space of geometric real functions to infer which functions are most representative of objects.

These methods are then in the end combined to give us the final, PHOG descriptor

Methods - Topological persistence

Topological persistence, or *persistence* for simplicity, is a theoretical framework well suited to shape analysis, comparison and matching.

It provides a qualitative, multi-scale description of a shape modeled by a pair (X, f), with X a topological space representing the object under study, and f :X→Rn a continuous function describing relevant shape properties of the object itself.

Topological persistence - 0th persistence diagram.

Points far from the diagonal describe important or global features, whereas points close to the diagonal describe local information, such as smaller details and noise.

The picture shows the 0th persistence diagram obtained from the 0th persistent Betti numbers of a surface model. The red vertical line can be seen as a point of infinity.

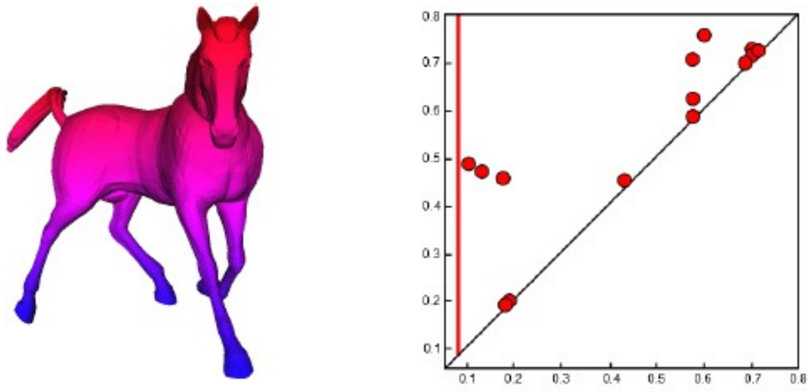

(a) (b)

(a) A model from the SHREC'13 benchmark with the "height" function color-coded (left), and the corresponding 0th persistence diagram (b).

Methods - Purely photometric description

Photometric properties can be represented in various different colour spaces, but the authors of this paper chose the CIELab colour space.

Some defining measures in CIELab colour space, is that, among other things, it has physiological support, meaning that it well represents the way the human eye sees colour, and that in this colour space, the tones and colours are held separately.

Methods - Purely photometric description

Integral geodesic functions when the weights ranges from purely geometric to purely photometric. The colour represents the function from low (blue), to high (red) values.

Methods - Hybrid Description

The main idea of the hybrid description is to define a geodesic distance in a higher dimensional embedding space, and use it as a real-valued function in the persistence framework.

The authors exploit the fact that the CIELab colour space is a perceptually uniform space, which implies that uniform changes of coordinates in the CIELab space corresponds to uniform changes in the colour perceived by the human eye. In other words, it is the best colour space for classifying colours in the same way a human would.

Methods - Geometric description

As geometry is captured through real functions, the authors of this paper have decided to, for each model, build a histogram to code the mutual relationships among a broad set of purely geometric functions defined on the same shape. Finally, they propose a manner to cluster these functions, and to simplify the description in order to better find similarities between the models.

Methods - Geometric description

The Mutual Distance Matrix, or MDM for short, is used to map a 3D object into a 2D matrix. This matrix has the properties that that the ith row/column of the MDM identifies all the distances of the ith function with respect to F, it provides a partial ordering of the distances between the function fi and the other ones.

Put simply, the minima of this histogram corresponds to functions that are qualitatively similar to the given one, while maxima highlight functions that significantly differ.

Methods - Geometric description

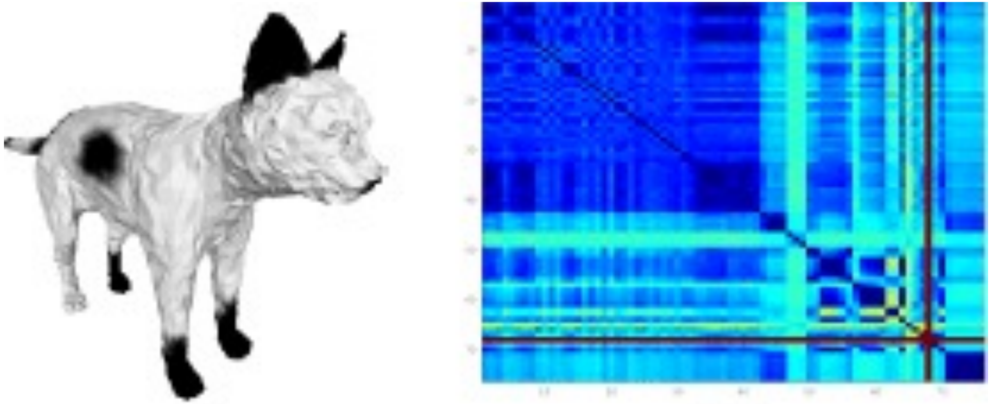

(a) (b)

(a) A model from the dataset and (b) the corresponding MDM matrix. The distances range from blue(zero) to red (1); large blue regions indicate that functions that are strongly similar.

Geometric description - Clustering

To cluster a set F of models, we can embed F via diffusion maps in a space where the Euclidian distance between two functions reflects their similarity.

To fulfill this demand, and at the same time make sure the distances among the cluster functions and the external ones are higher than the internal ones we can group the functions via diffusion map using the DBSCAN clustering technique. This technique is based on the notion of density reachability.

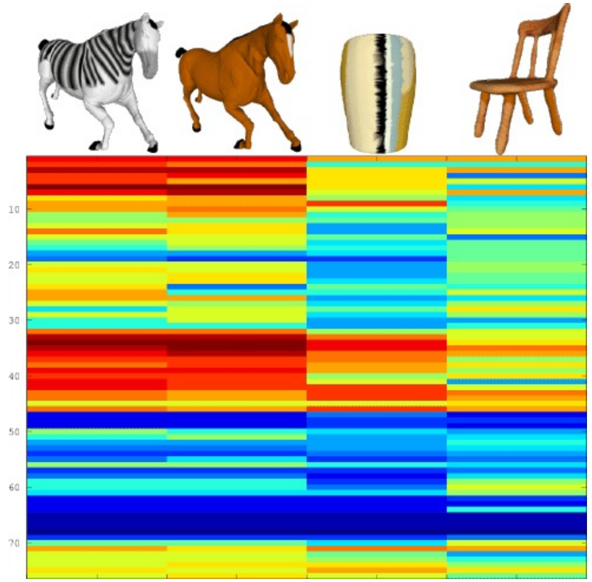

Results

To test the framework, the authors considered the SHREC' 13 data collection. This data set consists of 240 texture shapes in the form of triangle meshes.

Each triangle mesh is associated with their respective PHOG signature. It is a combination of the purely photometric, the hybrid, and the purely geometric descriptions. The final similarity score between two shapes is then obtained by summing up the three distances between the normalized signatures.

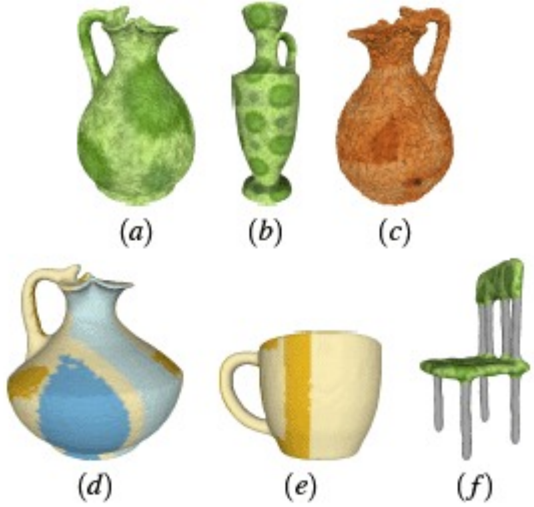

Some examples from the SHREC' 13 dataset used for testing

Results - A retrieval example comparison

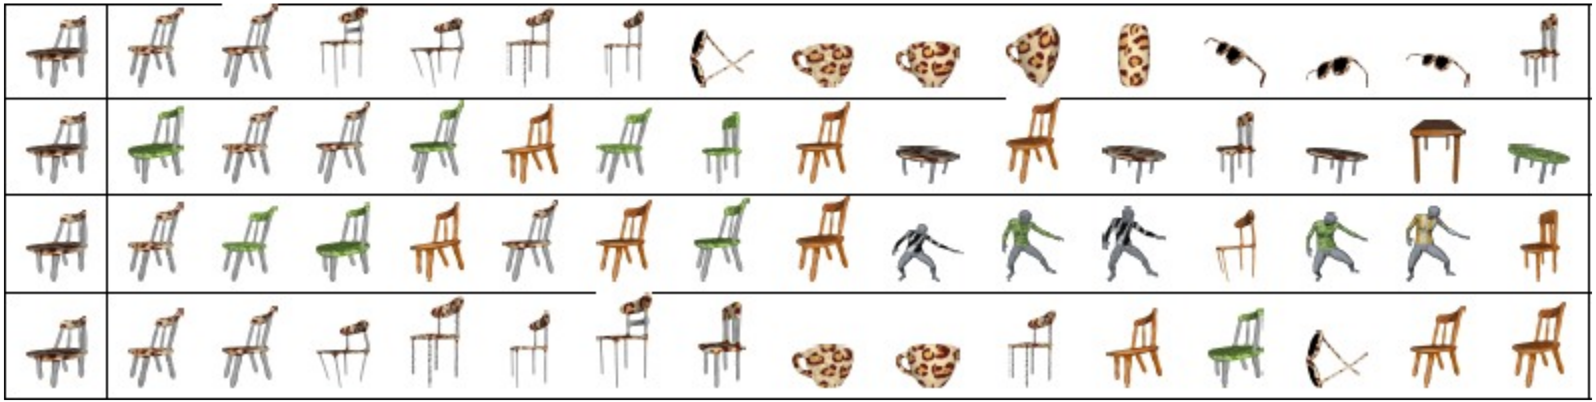

From top to bottom: retrieved items with respect to: LAB, hybrid, geometric, and combined description.

Results - Performance evaluation

The graph to the right shows the performance of all methods in terms of average precision-recall curves.
The larger the area below such a curve, the better the performance under examination.

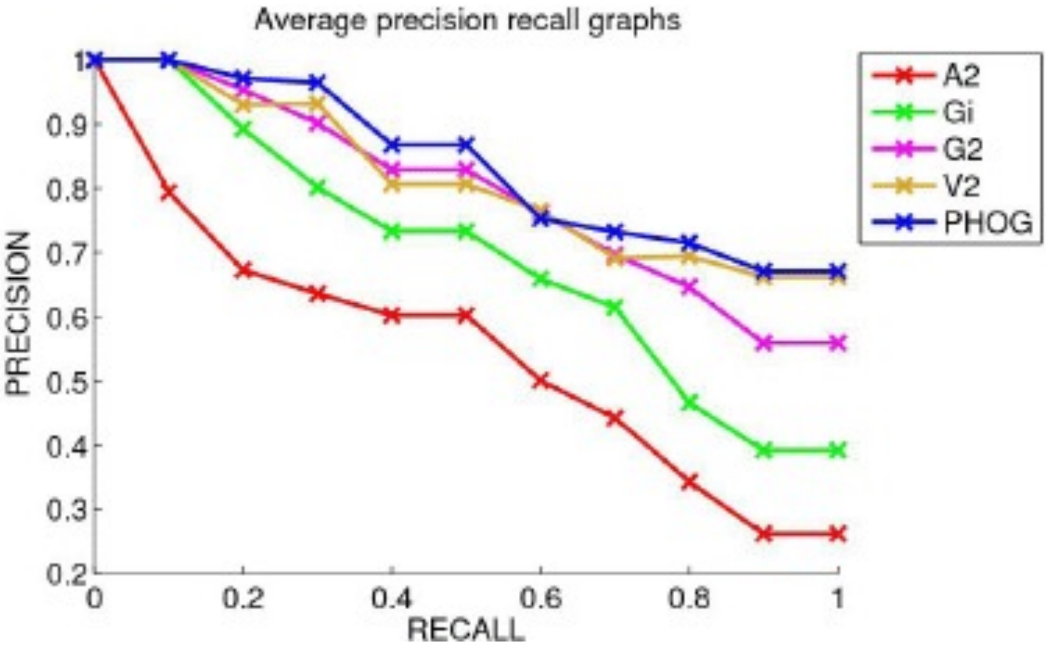

Precision-recall curves for all methods.

Results - Performance evaluation

The NDCG diagram to the right displays the performance evaluation for all methods according to the NDCG measure as a function of the rank P. NDCG is based on the assumption that highly relevant items are better if they have better ranks.

The higher the curve, the better the performance.

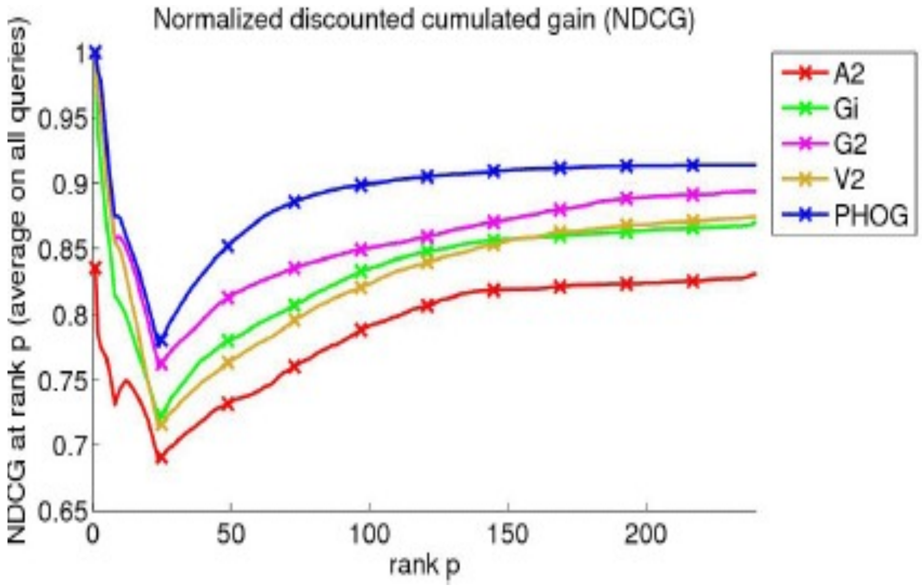

NDCG curves for all methods.

Conclusion

To target textured 3D object retrieval, the authors of this paper include photometric properties in the persistence framework. In that way they can inherit the invariance and robustness properties of the persistence setting, while also taking into account the texture of 3D objects. Then they fuse geometric, photometric and hybrid information to infer the similarity among objects.

This strategy offers several advantages. It is evolutionary (allowing other functions/properties to be included in the processing pipeline), modular (only informative information can be used) and adaptive (the way the similarity score is obtained may vary).

Conclusion

Improvements may be expected from the definition of new metrics, and thus new invariants, in the colour space.

The authors also predict that this technique could find a lot of application in domains where a texture is naturally associated to the 3D shape, such as cultural heritage, or face recognition, where the combination of photometric and geometric properties may increase trustworthiness under uncontrolled environmental conditions.

Questions