Eurographics Workshop on 3D Object Retrieval (2009)
I. Pratikakis, M. Spagnuolo, T. Theoharis, and R. Veltkamp (Editors)

Multilevel relevance feedback for 3D shape retrieval

D. Giorgi[1] and P. Frosini[2] and M. Spagnuolo[1] and B. Falcidieno[1]

[1]IMATI-CNR, Genova, Italy
[2]ARCES & Dept. of Mathematics, Università di Bologna, Italy

Abstract

Relevance feedback techniques are expected to play an important role in 3D search engines, as they help to bridge the semantic gap between the user and the system: similarity is a cognitive process, depending on the observer. We propose a novel relevance feedback technique, whose basic idea is threefold. First of all, the user is provided with a variety of shape descriptors, analysing different shape properties. The user then expresses her similarity concept through a friendly interface which supports multilevel relevance judgements. Finally, the system inhibits the role of the shape properties that do not reflect the user's idea of similarity. The assumption is that similarity may emerge as an inhibition of differences, i.e., as a lack of diversity with respect to the shape properties taken into account. The proposed technique is based on a simple scaling procedure, which does not require any a priori learning or optimization of parameters.

Categories and Subject Descriptors (according to ACM CCS): H.3.3 [Information Systems]: Information Search and Retrieval—Relevance feedback

1. Introduction

The quick increase in the number of available 3D models, along with the growing impact of 3D data in many areas such as medicine, bioinformatics, CAD and gaming [BKS*05], are creating the need for 3D search engines [TV08]: the retrieval of existing resources is indeed the basis for the reuse of the content and the embedded knowledge. In the IST Concertation Workshop "Challenges of Future Search Engines" (2005), it has been observed that the search of content "is an industry segment in its own right", whose "killer application is in everyday life". This is recognized for the traditional media, such as audio, images and video, whereas the impact of 3D in the societal and economic world makes us persuaded that this vision is going to come true for 3D shapes as well [Use08].

While designing a 3D search engine, we have to think that 3D shape retrieval is a complex interaction process between the user and the 3D content, along with its semantics. Smeulders et al. [SWS*00] put the attention on the *semantic gap*, i.e., the gap between the visual data information and the meaning of the data for the user. In specific domains (*narrow domains*), the semantic gap can be reduced by knowledge technologies (see the AIM@SHAPE approach [aim]). In a generic domain (*broad domain*), different strategies have to be adopted. First, we can develop *high level shape descriptors*, which code a variety of shape properties, so as to approximate the perceptual variety of humans. Then, we have to "include a human in the loop", that is, to make the user an active player in the search process [DJLW08]. As observed by Koenderink [Koe90], things possess a shape for the observer, in whose mind the association between the perception and the existing conceptual models takes place: similarity is a cognitive process, dependent on the observer.

The observer becomes part of the search process through the use of *relevance feedback* techniques [Roc71]. The latter allow the user to incorporate her judgement in a retrieval system by iterating three processes: the system returns a list of items in response to a query; the user gives a feedback about their relevance; the system updates the answers list, so that they better fit the user needs. Relevance feedback techniques help understanding the semantics of similarity for an observer, in the context of a specific query. Hence, they attempt to solve the semantic gap between description and meaning, between system and user.

In this paper, we devise a new relevance feedback technique to support 3D retrieval. The method takes advantage

DOI: 10.2312/3DOR/3DOR09/045-052

of state-of-the-art shape descriptors, which capture complementary shape properties, and helps the user to specify intutively, in a broad domain, the narrow shape domain she has in mind. Because of the intrinsic complexity and variability of 3D shapes, it is fundamental to go beyond a traditional, binary classification (relevant/not relevant). Hence, we define a technique based on multilevel relevance judgements expressed by the user through a friendly interface. We aim to approximate the similarity (pseudo)metric the user employs to compare objects: we start from the set of distances corresponding to shape descriptors, and progressively inhibit through a scaling procedure those distances that are not compatible with the user judgement.

The assumption is that similarity may emerge as an inhibition of differences, i.e., as a lack of diversity with respect to the shape properties taken into account. As suggested by the Longman Dictionary of Contemporary English, *shape* is "the outer form of something by which it can be seen (or felt) different by something else". This implies that it is an equivalence class of similarity that makes shape emerge from an object, whereas having a similar shape implies the absence of differences with respect to the shape properties which are relevant for the comparison (hence the user) at hand.

2. Related work

To the best of our knowledge, relevance feedback methods were introduced by Rocchio in 1971 for text document retrieval [Roc71]. Then, in the 1990s Rui et al. extended relevance feedback techniques to image retrieval [RWOM98]: see [ZT03] for a survey paper.

In the latest years, relevance feedback techniques have started attracting the attention of the 3D retrieval community [PPT*08]. The first proposal is by Elad et al. [ETA01], who use geometric moments for 3D retrieval and make the algorithm interactive by using Support Vector Machines. Bang and Chen [BC02] use feature vectors to describe 3D shapes, and propose a feedback technique based on a "morphing" strategy in the space of features: unlabeled objects are moved away or towards the query, according to their position with respect to labeled examples. Atmosukarto et al. [ALH05] propose to weight the ranking list of the retrieved objects, by pushing relevant and irrelevant objects to the top or the bottom of the list, respectively. Leifman et al. [LMT05] use Linear Discriminant Analysis and Biased Discriminant Analysis in a two-stage strategy, consisting of an off-line pre-processing step followed by the proper interactive phase. Recently, Onasoglou and Daras proposed two relevance feedback methods [OD08], which apply semantic forces of attraction or repulsion between 3D objects in a feature space. Semantic forces are built according to user judgements. In the first algorithm, they are combined with geometric forces determined by shape descriptors, while they stay purely semantic in the second algorithm.

Our algorithm proposes a continuous, numerical scale for relevance judgements, instead than a small, usually binary (relevant/not relevant) number of predefined relevance levels (see also [OD08]). Indeed, we aim to approximate the pseudometric the user employs when comparing 3D objects, according to the shape properties she is interested in. In addiction, our method combines multiple descriptors by working in the space of distances, thus being independent of the nature of the descriptors themselves. Finally, this technique does not require parameter optimization steps nor statistical learning procedures.

3. Relevance feedback technique: main ideas

This paper is devoted to propose a new method for the interactive approximation, based on the user feedback, of a pseudodistance δ on a dataset $\Sigma = \{x_1, \dots, x_N\}$. The pseudodistance δ represents the dissimilarities between the objects of Σ with respect to the subjective judgement of the user. A pseudodistance δ is a metrics without the property assuring that $\delta(x_i, x_j) = 0$ implies $x_i = x_j$ (in other words, two objects can have vanishing pseudodistance without coinciding).

We start with a family of 3D shape descriptors, producing a family $\mathcal{G} = \{d_1, \dots, d_n\}$ of pseudodistances between the objects in the database Σ. The way humans perceive and recognize things suggests that the recognition of any object requires a plurality of different recognitions, according to different object properties: we recognize a cat for having four legs, claws and a tail, for being furry, and so on. The same concept, in the context of 3D shape searching, turns into the need of describing and comparing 3D shapes according to different shape properties, i.e., according to different descriptors and comparison methodologies.

We define the initial pseudodistance between the objects in the database as the maximum amongst the distances encoding the variety of shape properties, $D = \max_{d \in \mathcal{G}} d$. Once a query is submitted, D is used to sort the database objects in order of decreasing similarity to the query. The system then returns a first list of answers. We use the max operator, instead of a traditional weighted linear combination of distances, because the sum operation is related to the "OR" operator, while the maximum is related to the "AND" operator, more suitable for the subjective comparison of complex shapes. In other words, we consider two objects similar if they are similar according to all the properties taken into account. The definition of the set of properties that are relevant for the user is the goal of the feedback technique.

We proceed by asking the user to give a feedback about the relevance of some answers. Due to the complexity of 3D shapes, and the variability of shape properties, we think that it is fundamental to go beyond a traditional, binary classification (relevant/not relevant), and let the user express the complexity of her judgement through a numerical scale [DJLW08, WLM03]. We provide the user with an interface endowed with a slider, so that she can move the cursor along the slider to express the similarity between two objects. The position of the cursor is turned into a numerical value, expressing the value of the pseudodistance δ. The latter is the (unknown) pseudodistance the user refers to for comparing the objects in the dataset Σ, and is the target of our approximation.

Given a query q, let us assume that the feedback given by the user implies the knowledge of the pseudodistance δ between q and the objects in a subset S of Σ, representing the user's opinion about the dissimilarities between the query and the objects in S. This knowledge is used to inhibit the role of the pseudodistances in the family $\mathcal{G}$ that are not compatible with the user's judgement, meaning that they perceive as different those objects that the user considers as similar. The mathematical idea is to rescale each pseudodistance d_i to a pseudodistance $\tilde{d_i} = \lambda d_i$, choosing λ equal to the larger value for which the equality $\lambda d_i(q,x_j) \leq \delta(q,x_j)$ holds when $x_j \in S$. In other words, we rescale d_i until it becomes compatible with the information given by the user, in the sense that the new pseudodistance associates each pair in $\{q\} \times S$ to a dissimilarity value not larger than the one expressed by the user, seen as the ground truth. This process produces a new family of pseudodistances $\widetilde{\mathcal{G}} = \{\tilde{d} : d \in \mathcal{G}\}$. After our rescaling procedure, we can use the new pseudodistance $\tilde{D} = \max_{\tilde{d} \in \widetilde{\mathcal{G}}} \tilde{d}$ for our retrieval purposes, as an approximation of δ. It is straightforward to verify that the resulting pseudodistance $\tilde{D}$ does not depend on the order in which the normalizations are performed.

The scaling procedure can be implemented as follows. Given a query object q, suppose that the user gives her numerical judgement $\delta(q,\bar{x})$ about the dissimilarity between the query and a returned object $\bar{x}$. This implies the substitution of each pseudodistance d_i with a new pseudodistance $\tilde{d_i}$ obtained by multiplying *everywhere* d_i by $\lambda_i = \min\{1; \frac{\delta(q,\bar{x})}{d_i(q,\bar{x})}\}$. This normalization procedure forces each d_i to respect the user judgement about the pair $(q,\bar{x})$: if $d_i(q,\bar{x}) \leq \delta(q,\bar{x})$, the pseudodistance d_i stays the same, while, if $d_i(q,\bar{x}) > \delta(q,\bar{x})$, d_i is lowered by multiplying everywhere its values, for each pair (q,x_k), $x_k \in \Sigma$, by the constant $\frac{\delta(q,\bar{x})}{d_i(q,\bar{x})} < 1$. This procedure corresponds to the cancellation of distances that are not compatible with the user judgement. Indeed, if the value of a pseudodistance d_i is much larger than a small dissimilarity perceived by the user, then the value λ_i becomes very small. Therefore, d_i is substituted with a pseudodistance that has very low values everywhere. This means that d_i plays very little role in the computation of the new pseudodistance $\tilde{D}$, being $\tilde{D}$ defined as $\tilde{D} = \max_{\tilde{d} \in \widetilde{\mathcal{G}}} \tilde{d}$.

The approach is based on three elementary but important properties of pseudodistances. Indeed, given the family $\mathcal{F}$ of all pseudodistance functions on a finite dataset Σ, the following statements hold :

1. For every subset $\mathcal{G} \subseteq \mathcal{F}$, the function D defined by setting $D(x_i,x_j) = \max_{d \in \mathcal{G}} d(x_i,x_j)$ is a pseudodistance on the dataset Σ;
2. For every non-negative $\lambda \in \mathbf{R}$ and every $d \in \mathcal{F}$ the function λd is a pseudodistance on Σ;
3. For every $d \in \mathcal{F}$ and every $x_i, x_j \in \Sigma$ the equality
$$d(x_i,x_j) = \max\{d'(x_i,x_j) : d' \in \mathcal{F}, d'(x_i,x_j) \leq d(x_i,x_j)\}$$
holds (in other words, each pseudodistance on Σ is the maximum of the pseudodistances that are everywhere less than d).

These statements are worth to be analyzed more carefully. They suggest the idea of approximating a pseudodistance δ by computing the maximum of a set of pseudodistances lower than δ. Indeed, the following result can be easily proved, referring to the previously defined sets $\mathcal{F}$ and Σ:

Assume that two subsets $\mathcal{G} \subseteq \mathcal{F}$, $T \subseteq \Sigma \times \Sigma$ and an $\varepsilon > 0$ are given, such that:

1. for all $d \in \mathcal{G}$ and for any $(x_r,x_s) \in T$, $d(x_r,x_s) \leq \delta(x_r,x_s)$;
2. for any $(x_i,x_j) \in T$ there exists $\overline{d} \in \mathcal{G}$ such that $\delta(x_i,x_j) - \overline{d}(x_i,x_j) \leq \varepsilon$.

Define $\tilde{D}(x_i,x_j) = \max_{d \in \mathcal{G}} d(x_i,x_j)$ for every $(x_i,x_j) \in T$. Then $\max_T \left|\delta - \tilde{D}\right| \leq \varepsilon$.

This simple statement assures that in presence of a rich enough family of pseudodistances, our method produces a good approximation of the pseudodistance δ representing our ground truth. In this framework, the normalization process can be seen as an easy way to enrich a finite family of pseudodistances. We observe that the definition of $\tilde{D}$ given in the previous statement is equivalent to the one previously described in this paper.

In our approach the pseudodistances are seen as judges of the dissimilarity, and similarity is interpreted as the absence of pseudodistances claiming the existence of dissimilarities. The learning procedure is based on the fact that each time we retrieve a false negative o_{FN} while looking for something similar to a queried object q, the relevance of the pseudodistances that erroneously see o_{FN} as quite different from q is reduced by the rescaling process. In other words, we can get rid of the pseudodistances that express wrong judgements by the rescaling process.

The advantage of our approach is threefold. First of all it is quite simple, and the rescaling process trivially converges to a pseudodistance D. On the other hand no optimization process is required, and the complexity of the algorithm (after computing the pseudodistances d_i depending on the chosen descriptors) is linear in the number of descriptors. Finally, the possibility of working in the space of distances between descriptors, rather than in the space of features, allows us to take advantage of heterogeneous shape descriptors, e.g., feature vectors and graphs.

A final remark is due about the use of a metrical approach to shape comparison. Indeed, this approach can be criticized by maintaining that shape comparison often does not obey the axioms of a metric. A classical example claims that a centaur can be similar both to a man and a horse, without the horse and the man being similar to each other. As a consequence, the triangular inequality could be considered inadequate for shape discrimination. This position follows from mixing two different pseudodistances d_H (expressing the dissimilarity with respect to the properties characterizing a horse) and d_M (expressing the dissimilarity with respect to the properties characterizing a man) by the min operator. Accepting this example corresponds to examining shape properties that depend not only on the user, but also on the pairs of objects that she is comparing. Indeed, it is based on the change of the shape properties the user evaluates, depending on what she is looking at. However, in this paper we are just interested in shape properties that are chosen by the user once and for all. For this kind of properties, the previously given example does not prove that a metrical approach to the problem is unfit, but just that a suitable selection and mixing procedure should be applied to our pseudodistances, similarly to what we do in this paper by the rescaling process and the max operator. In our setting, while d_H should be emphasized when the user is interested on the shape properties characterizing a man, d_M should be focused when the user is interested on the shape properties describing a horse. Mathematically speaking, d_H and d_M should not be mixed by considering the function $\min\{d_H, d_M\}$ since the min operator does not preserve the property of being a pseudodistance, while this property is preserved by the max operator.

4. Relevance feedback technique: implementation

In our prototypal implementation, the set of 3D shape descriptors include the spherical harmonics [KFR03] (implementation taken from [shw]); the lightfield descriptors [COTS03] (implementation taken from [lfw]); the size functions (our implementation) with the distance from the barycenter and the integral geodesic distance as measuring functions [FL99, BDF*]; the eigenvalues of the Laplace-Beltrami operator [Reu06] (our implementation); the descriptors participating in the Watertight Track of the SHape REtrieval Contest (SHREC) 2007 [Vt07], namely the multivariate density-based descriptor by Agkul et al. [ASSY06]; the depth line encoding descriptor by Chaouch et al. [CVB07]; the multi-view descriptor by Napoleon et al. [Vt07]; the spherical trace transform by Daras et al. [ZDA*07]; the augmented multiresolution Reeb graph by Tung and Schmitt [TS05]. By comparing these descriptors, we derive the set of initial distances $\mathcal{G} = \{d_1, \dots, d_{10}\}$. These distances are normalized so that their values are in the interval $[0,1]$: 0 stands for completely similar and 1 stands for completely different.

We used geometric, topological, structural and view-based descriptors, so as to capture a variety of different (preferably independent) 3D shape properties. Note that the comparison technique between the augmented multiresolution Reeb graphs does not induce a pseudometric (as required by our method) but a semimetric, i.e., the triangle inequality does not hold. In any case, we observed that in our database (the database of 300 models introduced in the Stability Track of SHREC 2008, see Section 5) the inequality is violated only in a very small number of cases (less than $10^{-5}\%$ of the triplets of objects). Hence, we included this descriptor as well.

The prototypal interface allows the user to select a query in the database, then returns a first list of items. The latter are ordered according to the increasing distances to the query, and shown in a HTML file. As explained, at this point the pseudodistance is $D = \max_{d\in\mathcal{G}} d$. Next, the user can select one or more objects amongst those returned by the system, and judge their similarity to the query. The multilevel judgement, corresponding to the user subjective pseudometric δ, is expressed through a slider: the position of the cursor indicates a value ranging in $[0,1]$. To help the user select the value, qualitative judgements (e.g., "very similar", "somewhat similar", "completely different", and so on) are associated with intervals of values and printed next to the slider, as the cursor moves. The slider is used to provide the user with an intuitive and friendly tool to express a multilevel relevance judgement, ranging in a (quantized) continuous interval. This is the main difference with respect to other relevance feedback methods for 3D retrieval, which only allow the user to select amongst a small number of relevance levels, usually two (relevant/not relevant). With our technique, the user is left free to express her feelings through a continuous scale, which reflects her internal idea of distances between objects. At the same time, the method is robust with respect to small changes in the value of the user judgement, as a small change in the value of $\delta(q,\bar{x})$ induces a small change in the multiplying factor $\frac{\delta(q,\bar{x})}{d_i(q,\bar{x})}$. Consequently, a drastic change in the scaling procedure only results from big changes in the user judgements, that are not expected to happen neither by chance or by mistake.

Our prototypal interface is shown in Figure 1. A window allows the user to choose a query in the database, and select and judge through the slider the items returned by the system. The prototype is implemented using the C++ language under the Microsoft Visual Studio 2005 environment.

5. Experiments and discussions

We conducted a series of experiments using the dataset of 300 models used in the Stability Track of the SHREC Contest 2008, so as to evaluate the potentialities of our relevance feedback technique. The dataset is made of 15 classes, with 20 models per class. The objects range from humans and animals to cups and mechanical parts: see the Track Report for further details [AB08].

Figure 2 shows how our method is able to refine a query,

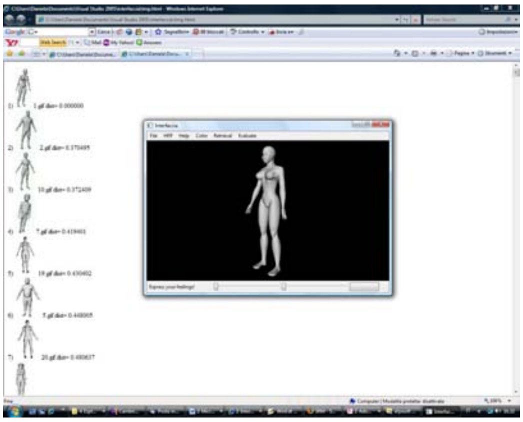

Figure 1: *The prototypal interface.*

in terms of improved precision and recall. In this example, we are looking for a member of the "armadillo" class, shown on the left. On top right, we find the list of the first returned items (apart from the 1st one, which is always the query itself and is not shown) ordered according to D, i.e., sorted by the maximum value taken by the distances d_is. The results, though not perfect, are nice, meaning that most of the d_is are able to recognize the armadillos. Some false positives may be observed, namely four horses. With a first feedback iteration, we communicate to the system that we are looking for armadillos like the one in the blue circle: we move the cursor along the slider and set the dissimilarity value δ between the query and the circled armadillo to 0.25. Figure 2 (bottom) shows the updated list. With a single feedback iteration, we managed to get rid of the unwanted horses, and get three new relevant armadillos, marked with green rectangles.

Let us see what happens when dealing with a more complicated class, showing higher shape variability. In Figure 3, we query the system using a spring. The results list contains little springs and many false positives (Figure 3 (top)). We label the spring enclosed in the blue circle with a value of δ equal to 0.15, to suggest that we are looking for a differently curved type of springs. Figure 3 (bottom) shows the list updated after the feedback iteration, which is populated by many other springs. This happens because the system automatically inhibits the distances which consider only metric properties of springs. Figure 4 shows a similar example.

As seen, good results can be obtained after a single iteration, possibly by judging a single item. This is very important, because a relevance feedback technique is useful as long as it is not boring for the user, i.e., it does not require too many iterations and relevance judgements. We also remark that, though examples are not shown for space reasons, if we slightly change the values of the user subjective judgement, namely the values of δ, then the system answers stay almost the same, as discussed in Section 4.

The next examples we are about to discuss rise some questions and suggest some directions for future research. In Figure 5 (top), the system gets confused between humans and armadillos. In the first feedback iteration, we move the cursor along the slider and mark with 0.2 the walking woman in the blue circle. The system selects the descriptors that agree with our idea. Consequently, many other humans appear in the list (Figure 5 (middle)). Suppose now that we are looking for a particular kind of human model, namely a sitting woman. Hence, in a second iteration, we feed the system with a value equal to 0.10 for the distance between the query and the sitting woman in the blue circle. What happens (see Figure 5 (bottom)) is that we manage to find another sitting woman (marked with a green rectangle), but the list now also includes a false positive, namely the four legs animal marked with the red rectangle. We believe this depends on the fact that the descriptors we are using do not refer to "atomical" properties: there is no descriptor looking at a single property (e.g., having legs or hands, being seated and so on). This may cause the presence of false positives.

The choice of descriptors deserves some attention. Our method is able to produce results according to some wanted shape properties as long as the properties are represented by the descriptors involved. In other words, we start with a set of descriptors and attempt, with our feedback technique, to select, in a query-dependent yet automatic manner, the right subset of descriptors capturing the properties the user is interested in. The point is, we can only do that if the system includes some descriptors perceiving these properties. Ideally, for each shape property in the user mind, there should be a descriptor devoted to capture it, and possiblly only it. After various feedback iterations, the user would be given the objects with the sought properties. This is unfeasible in practice, because of the limits of the current state of the art, and also because the number of descriptors must respect a tradeoff between efficiency and description efficacy. A promising direction is given by the study of the descriptors that are parametric with respect to different shape properties [BFF*07], in particular the descriptors originating in Morse theory, such as Reeb graphs and size functions [BDF*]. Moreover, we believe the community should welcome a deep study of the relationship between shape properties and shape descriptors: current studies mainly focus on a quantitative assessment of descriptors [DP06], in terms of their performance on existing benchmarks [Vt07], whereas the literature confirms that the performance is query dependent. New and extensive studies focusing on a qualitative evaluation of the descriptors are needed, if we want to build dynamic, user-centric search engines.

The example in Figure 6 poses an important question. As the user searches for a small table with three legs, the system does not return any similar table amongst the first retrieved models, which are all false positives (Figure 6 (top)). In this case, the user is not able to feed the system with examples of objects sharing the shape properties she has in mind. Obvi-

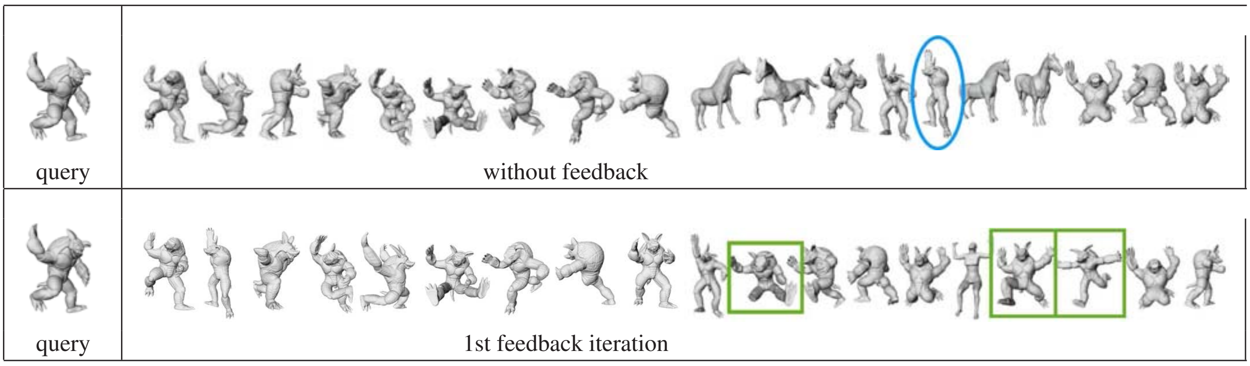

Figure 2: *Retrieval session for an armadillo model. The user judges the circled armadillo to be at distance 0.25 from the query. The system updates its answer and shows the three new armadillos in the green rectangles. With a single feedback iteration, precision improves from 0.79 to 0.95.*

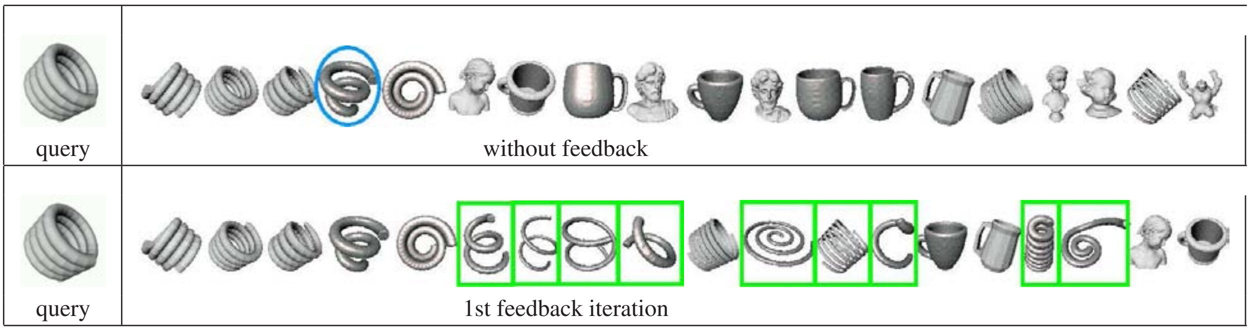

Figure 3: *Retrieval session for a spring. The user judges the circled spring to be at distance 0.15 from the query. The system updates its answer, showing a number of different springs. With a single iteration, recall improves from 0.32 to 0.8.*

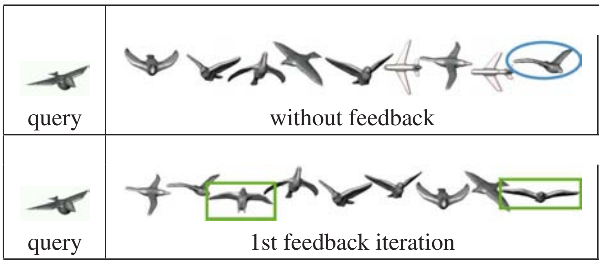

Figure 4: *Retrieval session for a bird model.*

ously, it is not practical to ask the user to scroll the list until she finds a good example. A first idea is to look for positive examples in lists of answers responding to different criteria. A second option is to ask the user to evaluate negative examples, i.e., to give negative feedback. Both the problems of maximizing the number of positive examples and of treating negative examples are open problems in the literature.

According to the first idea, to solve a pathological situation like that in Figure 6 (top), in which the initial distance *D* fails to produce good results, we could show the user a different set of models to search in, selected according to a different criterion than *D*. We propose to show the user the set of models given by the first returned item (apart from the query itself) for each distance, under the assumption that the shape properties the user is looking for are represented by some of the descriptors. Computing this “melting pot” of models is more practical than asking the user to look into the lists of results of each descriptor. Following this idea, in the example in Figure 6, we manage to find a pair of tables to be used as examples (Figure 6 (middle)), identified by the size functions with the integral geodesic distance and the augmented multiresolution Reeb graph, respectively. If the user judges, through the interface slider, the distance between the query and the table in the blue circle to be 0.1, the system updates the outcome and returns many other tables (Figure 6 (bottom)). Although this idea seems to work in many situations, more research is needed.

As for the treatment of negative examples, some proposals have been presented in the literature about relevance feedback [KZB06], whereas there is a debate about negative examples having destructive effects [DJLW08]. With our as-

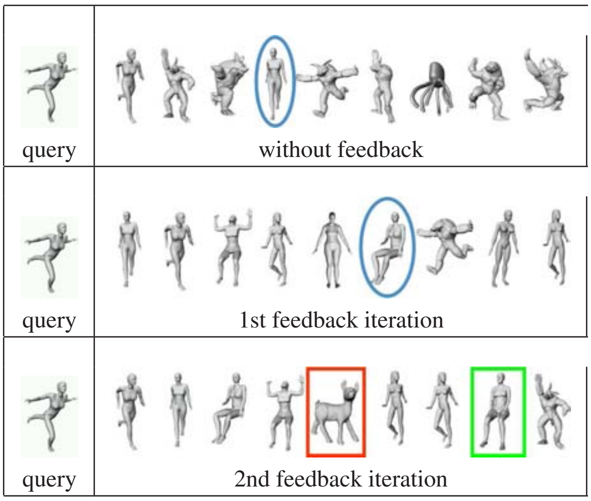

Figure 5: *Retrieval session for a woman model. Top: the user gives a feedback of* $\delta = 0.2$ *for the walking model. Middle: the system updates its answer. The user refines her query, by selecting the sitting model and suggesting* $\delta = 0.1$. *Bottom: the results obtained at the second iteration: a new sitting woman appears, along with a false positive.*

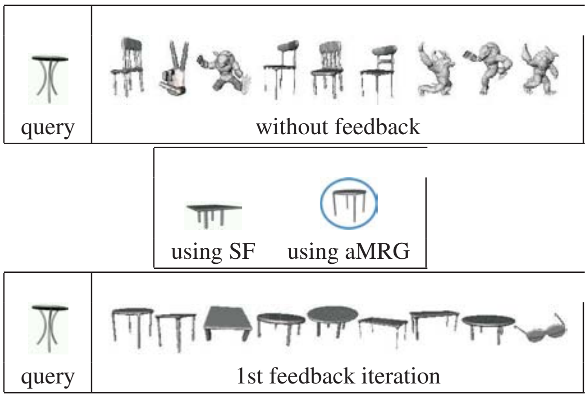

Figure 6: *Retrieval session for a table model. Top: the first outcome of the system does not contain any similar table. Middle: two tables are found in a different set of answers; the user gives her feedback, by stating that* $\delta = 0.1$ *is the distance between the query and the circled table. Bottom: after the feedback step, the updated list of models shows a set of tables to the user, possibly answering to her query.*

sumptions, negative examples are not useful, because our method works by inhibiting those distances that perceive as different the objects that the user considers as similar: objects judged to be not similar give little information to the system, as they do not communicate what the user is looking for, but only what she is *not* looking for. We are currently investigating the possibility of extending the method to take advantage of negative examples.

6. Conclusions

In this paper, we have presented a novel relevance feedback algorithm for 3D shape retrieval, whose aim is to approximate the pseudometric employed by a user to compare objects. The algorithm is based on a simple scaling procedure, and motivated by elementary yet important mathematical properties of pseudodistances.

We have implemented a prototype interface, enabling the user to express her relevance judgement through a gauge of values. Note that, although many research works investigated appropriate ways to measure the degree of users' relevance judgments, there is still no consensus regarding how to design such a measurement scale [WLM02]. This topic deserves further research.

We are currently organizing a more complete series of experiments, which require the active collaboration of a set of volunteer users. Since traditional performance measures such as precision and recall are based on dichotomy relevance judgment, we are considering the use of other measures which can deal with gauges of values [Yao95]. Also, the evolution of the retrieval performance with the number of user feedback iterations deserves a deeper investigation.

Further directions of future research include the study of all complementary steps that decide the success of a relevance feedback method: facilities for query selection and database navigation, aimed to solve the "page zero problem"; a display space that better presents the outcome to the user; personalized mechanisms based on the user's profile and needs [SWS*00].

Acknowledgements

Work supported by the Project SHALOM funded by the Italian Ministry of Research (contract number RBIN04HWR8) and the EU FP7 Project FOCUS K3D (contract number ICT-2007.4.2). Work partially developed within the activity of ARCES "E. De Castro", University of Bologna, under the auspices of INdAM-GNSAGA.

References

[AB08] ATTENE M., BIASOTTI S.: Shape retrieval contest 2008: Stability of watertight models. In *Proc. IEEE SMI 2008* (2008), pp. 219–220.

[aim] http://aimatshape.net.

[ALH05] ATMOSUKARTO I., LEOW W., HUANG Z.: Feature combination and relevance feedback for 3D model retrieval. In *Proc. 11th Multimedia Modelling Conference* (2005), pp. 334–339.

[ASSY06] AKGÜL C., SANKUR B., SCHMITT F., YEMEZ Y.: Density-based shape descriptors for 3D object retrieval. In *Proc. MRCS* (2006), vol. LNCS 4105, pp. 322–329.

[BC02] BANG H., CHEN T.: Feature space warping: an approach to relevance feedback. In *Proc. IEEE Int. Conf. on Image Processing* (2002), pp. 22–25.

[BDF*] BIASOTTI S., DE FLORIANI L., FALCIDIENO B., FROSINI P., GIORGI D., LANDI C., PAPALEO L., SPAGNUOLO M.: Describing shapes by geometrical-topological properties of real functions. *ACM Computing Surveys 40*, 4, 12:1–12:87.

[BFF*07] BIASOTTI S., FALCIDIENO B., FROSINI P., GIORGI D., LANDI C., MARINI S., PATANÉ G., SPAGNUOLO M.: 3d shape description and matching based on properties of real functions. In *Eurographics 2007 Tutorial Notes* (2007), The Eurographics Association, pp. 1025–1074.

[BKS*05] BUSTOS B., KEIM D., SAUPE D., SCHREK T., VRANIC D.: Feature-based similarity search in 3D object databases. *ACM CSUR 37*, 4 (2005), 345–387.

[COTS03] CHEN D., OUHYOUNG M., TIAN X., SHEN Y.: On visual similarity based 3D model retrieval. *Computer Graphics and Applications*, 22 (2003), 223–232.

[CVB07] CHAOUCH M., VERROUST-BLONDET A.: 3d model retrieval based on depth line descriptor. In *Proc. ICME'07* (2007), IEEE.

[DJLW08] DATTA R., JOSHI D., LI J., WANG J.: Image retrieval: ideas, influences and trends of the new age. *ACM Computing Surveys 40*, 2 (2008).

[DP06] DEL BIMBO A., PALA P.: Content-based retrieval of 3D models. *ACM Tran Multimedia Computing, Communications and Applications 2*, 1 (2006), 20–43.

[ETA01] ELAD M., TAL A., AR S.: Content-based retrieval of VRML objects - an iterative and interactive approach. In *Proc. EG Multimedia* (2001), pp. 97–108.

[FL99] FROSINI P., LANDI C.: Size theory as a topological tool for Computer Vision. *Pattern Recognition and Image Analysis 9* (1999), 596–603.

[KFR03] KAZHDAN M., FUNKHOUSER T., RUSINKIEWICZ S.: Rotation invariant spherical harmonic representation of 3D shape descriptors. In *Proceedings of SGP 2003: Eurographics Symposium on Geometry Processing* (2003), pp. 156–165.

[Koe90] KOENDERINK J.: *Solid shape*. Cambridge, MA: MIT Press, 1990.

[KZB06] KHERFI M., ZIOU D., BERNARDI A.: Combining positive and negative examples in relevance feedback for content-based image retrieval. *J. Vis. Commun. Image R. 14*, 3 (2006), 428–547.

[lfw] http://3d.csie.ntu.edu.tw/dynamic/3dretrieval/index.html.

[LMT05] LEIFMAN G., MEIR R., TAL A.: Semantic-oriented 3D shape retrieval using relevance feedback. *The Visual Computer 21* (2005), 865–875.

[OD08] ONASOGLOU E., DARAS P.: Semantic force relevance feedback, content free 3d object retrieval and annotation propagation: bridging the gap and beyond. *Multimed Tools Appl 39* (2008), 217–241.

[PPT*08] PAPADAKIS P., PRATIKAKIS I., TRAFALIS T., THEOHARIS T., PERANTONIS S.: Relevance feedback in content-based 3D object retrieval: A comparative study. *CAD and Applications 5*, 5 (2008), 753–763.

[Reu06] REUTER M.: *Laplace Spectra for Shape Recognition*. Books on Demand, ISBN 3-8334-5071-1, 2006.

[Roc71] ROCCHIO J.: *Relevance feedback in Information Retrieval*. Prentice Hall, 1971.

[RWOM98] RUI Y., WANG T., ORTEGA M., METHOTRA S.: Relevance feedback: a power tool in interactive content-based image retrieval. *IEEE Tran. Circuits and Systems for Video Tech. 8*, 5 (1998), 644–655.

[shw] http://www.cs.jhu.edu/misha/.

[SWS*00] SMEULDERS A., WORRING M., SANTINI S., GUPTA A., JAIN R.: Content-based image retrieval at the end of the early years. *IEEE Trans. PAMI 22*, 12 (2000).

[TS05] TUNG T., SCHMITT F.: The augmented multiresolution Reeb graph approach for content-based retrieval of 3D shapes. *IJSM 11*, 1 (2005).

[TV08] TANGELDER J., VELTKAMP R.: A survey of content-based 3D shape retrieval methods. *Multimedia Tools and Applications 39* (2008).

[Use08] USER CENTRIC MEDIA CLUSTER OF EU PROJECTS, POSITION PAPER: Research on Future Media Internet in a global context.

[Vt07] VELTKAMP R., TER HAAR F.: *3D Shape Retrieval Contest*. Tech. Rep. UU-CS-2007-015, 2007.

[WLM02] WU H., LU H., MA S.: Willhunter: Interactive image retrieval with multilevel relevance measurement. In *Proc. IEEE Int. Conf. on Image Processing* (2002), pp. 22–25.

[WLM03] WU H., LU H., MA S.: Multilevel relevance judgement, loss function and performance measure in image retrieval. In *Proc. CIVR* (2003), vol. LNCS 2728.

[Yao95] YAO Y.: Measuring retrieval effectiveness based on user preference of documents. *Journal of the American Society for Information Science 39*, 2 (1995), 133–145.

[ZDA*07] ZARPALAS D., DARAS P., AXENOPOULOS A., TZOVARAS D., STRINTZIS M.: 3d model search and retrieval using the spherical trace transform. *EURASIP Journal on Advances in Signal Processing* (2007).

[ZT03] ZHOU X., THOMAS S.: Relevance feedback in image retrieval: a comprehensive review. *Multimedia Syst 5*, 8 (2003), 536–544.