

Utilizing treemaps for multicriterial search of 3D objects

Georgios Petkos, Vasilios Darlagiannis, Konstantinos Moustakas and Dimitrios Tzovaras

Center for Research and Technology Hellas, Informatics and Telematics Institute, Greece

Abstract

We propose a treemap based interface for presentation of search results according to multiple search criteria. Different colors are used to represent the relevance of each item in the database according to the different search criteria, while at the same time the treemap based representation allows the user to visually identify relevant groups of data, exploiting the hierarchical organization of the items in the database. Items are ranked according to each criterion and an aggregate ranking is computed using the Borda algorithm. Furthermore, appropriate interaction mechanisms are provided in order to assist the user in refining the presentation of the returned items and weigh the contribution of different criteria for retrieving combined search results. The Princeton benchmark 3D object database is used for this study, nevertheless the technique presented is appropriate for any multicriterial search application and in particular in cases where the data is organized hierarchically.

1. Introduction

Interesting work in the fields of visual analytics and information visualization for rich representation of search results and advanced interactive mechanisms has recently appeared. An example is ResultMaps [CDF09], which represents the hierarchical structure of search results using treemaps. [GPL*09] presents a tool that can cluster images returned from Google's image search and interact with the user to determine the most relevant search results. Similar applications can be found in [FGL*08] and [FKG*09], where the developed tools provide visual summarization of a large number of images retrieved from Flickr, in order to assist search and to provide helpful recommendations to the user.

This paper presents an interface that provides the user with a compact representation of the relevance of search results according to *multiple criteria* simultaneously. The developed interface utilizes a treemap in order to visualize the hierarchical structure of the items in the database, enabling the user to visually identify relevant groups of data for further inspection. Appropriately chosen colors in an appropriately chosen layout are used, so as to represent the relevance of the items according to the set of criteria in a non-cluttered manner. In addition, interaction mechanisms for managing the display of results and for refining the search are developed.

The focus of this study is on 3D object retrieval, however the methods presented are applicable to any other search application where the inclusion of multiple criteria is relevant. 3D object retrieval [BKS*06] [BKSS07] has become an important research area with many applications in fields like medicine, manufacturing, computer aided design and molecular biology. 3D object retrieval is a sort of content-based retrieval, which is in some sense a fuzzy operation. Compared to retrieval in relational databases, its goal is to identify items that cannot match exactly but are somehow similar to some query item or criteria. The common approach to content-based retrieval is to extract a set of features for each item in the database and define appropriate similarity measures based on these features. Given a search item, the features are extracted from it and the similarity measure to the items in the database is computed. Items are then ranked based on this measure and the most similar ones are returned to the user. In a slightly different scenario, a set of search criteria or features may be directly provided instead of being extracted from a search item.

The rest of the paper is organized as follows. A short discussion on 3D object features and the extraction process is first provided. Then, the ranking aggregation method that is used to fuse the rankings obtained by the set of criteria is outlined. Subsequently, the developed interface, i.e. the visual representation and interaction techniques, is described. A user evaluation study follows and finally, the paper con-

cludes with a discussion about the merits of the presented method and future improvements.

2. 3D object feature extraction

There is a multitude of methods in the literature for extracting geometric features with rich discriminative power. For instance, inspired by relevant work in image processing, [MAD*06] discusses the use of weighted Krawtchouk moments, [ZDA*07] uses a spherical trace transform and [MDTS09] uses an impact descriptor.

In this study, a single geometric descriptor, hybrid Krawtchouk moments and a set of soft shape features, like the volume, the ratio of volume to the volume of the convex hull and the surface of a 3D item are used as search criteria. We call them soft features as, in general, themselves are not sufficient for identifying items that belong in the same class in a very strict sense. Nevertheless, we can expect that specific values of these features or combinations of values are useful for identifying and therefore retrieving a class of objects. For instance, a cup and a ball may have the same volume, they may be distinguished though by the ratio of their volume to the corresponding convex hull. On the other hand, geometric descriptors are much more powerful for characterizing the similarity of 3D items. However, they are generally not intuitive, like the soft features are.

In the following, we outline the process of extraction of Krawtchouk moments. Krawtchouk moments are based on a set of orthonormal polynomials, associated with the binomial distribution, though more recent approaches expressed Krawtchouk polynomials in terms of hypergeometric functions. These fundamental concepts have been applied in 2D image analysis [YPO03], which have been extended to fit for 3D objects in [MAD*06] by introducing discrete weighted 3D Krawtchouk moments. The latter work has formed the basis for the experimentation of the proposed method. More particularly, the weighted 3D Krawtchouk moments of order $(n + m + l)$ of f are described as follows:

$$\begin{aligned} \bar{Q}_{nml} &= \sum_{x=0}^{N-1} \sum_{y=0}^{M-1} \sum_{z=0}^{L-1} \bar{K}_n(x; p_x, N-1) \\ &\times \bar{K}_m(y; p_y, M-1) \times \bar{K}_l(z; p_z, L-1) \times f(x, y, z) \end{aligned}$$

where $K_n(x; p, N)$ is the n -order Krawtchouk classical polynomials defined in terms of a hyper-geometric function. Weighted 3D Krawtchouk moments can be used as a descriptor of any 3D object, if it can be described by a function $f(x, y, z)$ defined in a discrete space $[0 \dots N-1][0 \dots M-1][0 \dots L-1]$, as long as the model is expressed as a binary volumetric function. However, a 3D model M is generally described by a 3D mesh. Therefore, in order to compute the weighted 3D Krawtchouk moments, the 3D mesh representation has to be converted into a volumetric representation

using an appropriate voxelization method as presented below:

The descriptor vector is composed of Weighted 3D Krawtchouk moments up to order s :

$$D = [\bar{Q}_{nml} | n + m + l \in [0 \dots s]]$$

In our implementation we combine the Weighted 3D Krawtchouk moments with the Spherical Trace Transform using 2D Krawtchouk, resulting in feature vectors with 37 discriminative elements, which in some of our previous work have shown remarkable results.

The soft features are computed using a voxelized representation of the 3D models of the items.

3. Ranking aggregation

The similarity of items according to different criteria needs to be merged in some way, so that an aggregate ranking can be obtained. A possibility would be to use probabilistic relevance methods [JWR00]. This would however require the definition of probability distributions of the features given that they are classified as relevant or not. Instead, we choose to merge the results at the ranking level and use the popular Borda method [SvZ09] [CFR06]. This works as follows.

We define a ranking π of N elements as a mapping $\pi : (1 \dots N) \rightarrow (1 \dots N)$. Given a set of K rankings $[\pi_1 \dots \pi_K]$, the Borda count is

$$Borda(i) = \frac{1}{K} \sum_{k=1}^K \pi_k(i).$$

The final ranking is obtained by sorting the N Borda counts.

This allows to obtain very fast a ranking of the relevance of according to multiple criteria. Nevertheless, it would be useful to allow the user to define how important different criteria are for searching 3D items. The user may even decide that some features are not important for 3D item search. For instance, given a query point, the user may decide that only the geometric features are important for determining the relevance of items to some query item and decide to leave the others out. In another scenario, the user may believe that the relevance of some item according to its volume is more important than the relevance according to surface. To this purpose, we use weighted Borda counts as:

$$Borda(i) = \frac{1}{K} \sum_{k=1}^K w_k \pi_k(i)$$

where the weights w_k sum to 1 and denote the relevant importance of each ranking.

4. Visual representation and interactive search

Provided that information on the hierarchical organization of the database is available, it is useful to visualize this information in order to assist the user in gaining a better insight

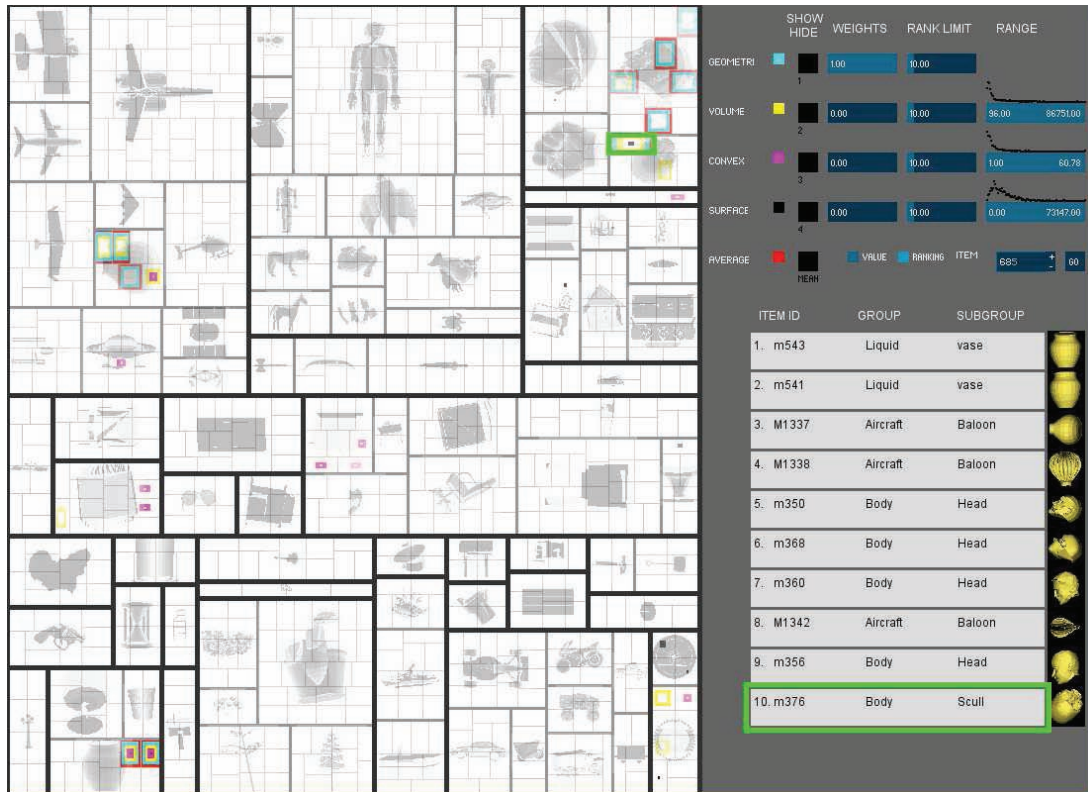


Figure 1: The complete search interface. A search by item is performed and the 10 top items in the ranking are displayed. It can clearly be seen that two items that belong to the subgroup "vase" are the most relevant to the query based on all features. Furthermore, it appears that two other subgroups of items are highly relevant to the query at least according to the geometrical descriptors.

on the results. There are many visualization methods that are suitable for representing the hierarchical organization of data, e.g. cladograms, hyperbolic trees and conetrees. Nevertheless, treemaps [JS91] tend to exploit the space they are allocated better than other methods and are therefore chosen for this study.

The search interface is built using the Processing environment. The squarified layout [BHvW99] was chosen for the treemap so that each item is represented using an as much as possible uniform shape. The Princeton benchmark database, which consists of 905 objects in 35 categories and 92 sub-categories is used. The search interface and treemap can be seen in Fig. 1. Groups and subgroups of items are indicated by thicker black and grey lines respectively.

A useful feature of the treemap is that it displays a 2D view of the most relevant to the current query item from each subgroup of items, in the treemap itself, thereby improving the perceptual processing of results by the user. It is important to stress that the displayed item changes dynamically according to the query, instead of just being static, display-

ing the median object for instance, providing a better clue about the relevance of each group of items.

Apart from the geometric descriptor that has been described, the aforementioned soft features, i.e. the volume, the ratio of volume to the volume of the corresponding convex hull and the surface of objects are used as search criteria. There are two options for search. The first is by value, where range bars are used to filter the displayed results according to the values that the user has provided. The distributions of the relevant features are displayed above the range bars to provide a hint to the user about the amount of items that will be retrieved. The other option is search by item, where the features are automatically extracted from the query item and are used to rank the items in the database according to their similarity to the features. When searching by item, the user can choose how many top ranked results will be displayed by manipulating the rank limit bar for each of the features.

As mentioned, different colors indicate the relevance / value of different features, i.e. the intensity of the color is proportional to the value or the ranking depending on the

mode of search. For instance, when searching by value, the user can determine a range of values for each feature that he is interested in. Items with larger feature values are displayed with higher intensity colors whereas items with lower values are displayed with lower intensity colors. Items with feature values that are out of the specified range are coded with white. Similarly, when searching by item, the items that have feature values that are closest to the feature value of the query item are displayed with higher intensity colors for the corresponding features. The user can determine which features are displayed so that he can spot relevant items according to the set of features that he deems important. Different arrangements like horizontal and vertical strips and adjacent rectangles have been tried, they were found however to create confusion about which features belong to each item and were therefore abandoned. Instead, concentric rectangles were found to maintain the cohesiveness of the visual representation of each item and were therefore chosen. An appropriate choice of ranges of colors for each feature has ensured that different items are easily compared regardless of what features the user has chosen to display. Relevant items can be easily spotted using this representation if they are characterized by all the colors that correspond to the criteria that the user has determined as visible. Moreover, similar groups of items in terms of some feature may be identified by regions where this feature is present. Additionally, the user can determine the relevant weights w_k of the ranking of each criterion for determining the aggregate ranking using Borda's algorithm. This enables him to simultaneously see the relevance according to different criteria and at the same time control the aggregation for assisting him/her in retrieving the most important results.

In addition, a list displays the 10 most relevant results, according to the aggregated ranking. Apart from the name of the item, the name of its group and subgroup are also displayed. This list is connected to the treemap and when an item in the list is highlighted, the corresponding item in the treemap is also highlighted and vice versa. If a highlighted item in the treemap is not among the top 10 in ranking, its details are displayed separately. The list is an important feature as it provides subtle details that cannot be easily incorporated in the visualization. For instance, it is difficult to disambiguate ranking between closely ranked items solely based on the treemap. Moreover, it provides more details about the groupwise coherence of the top results.

Figs. 1-3 describe a simple search scenario and show some of these features. Fig. 1 shows the results for a search by item. We assume that we are initially interested only in the most relevant items, so we only display the 10 top ranked items for each criterion. It can easily be seen that two items near the bottom left of the treemap, that belong to the subgroup "vase", appear to be relevant to the query according to all of the criteria. It can also be observed that two groups of items seem highly relevant to the query judging by the geometric similarity of some of the items belonging to these




Figure 2: Results for the same query as Fig. 1 but with the 50 top ranked items displayed. Apart from the already identified most relevant items on the bottom left, some other groups of similar items can easily be spotted. In particular, two clusters of items that are similar to the query item primarily in terms of the geometric descriptors (denoted by the cyan color) but some of them in terms of the other features as well.

groups to the query item. In order to verify this, we increase the number of the top ranked results that are returned. As it can be seen in Fig. 2, even more items from these groups are returned. In particular, the group at the top right and the group near the top left, seem highly relevant at least in terms of the geometric descriptors (denoted by the cyan color) but also in terms of the other features for some of them. Inspecting these groups of items using the list, one can see that they contain "balloons" and "heads", which may indeed be geometrically similar to some kind of "vase". Then suppose that the user decided that out of these items, the ones that are also similar in terms of volume are interesting for further inspection, the user has two options: either to decrease the weight of the other features and increase the weight of the volume feature or to hide some of the features that are not critical to the user search. This is displayed in Fig. 3, where the same results as in Fig. 2 are displayed, however the relevance according to the ratio of the volume of the item to the volume of the corresponding convex hull and the surface are omitted from the visualization. This provides a more clear view of the relevance of the items according to the criteria of interest, in this case the geometry and the volume. Consequently, individual items that are similar to the query item

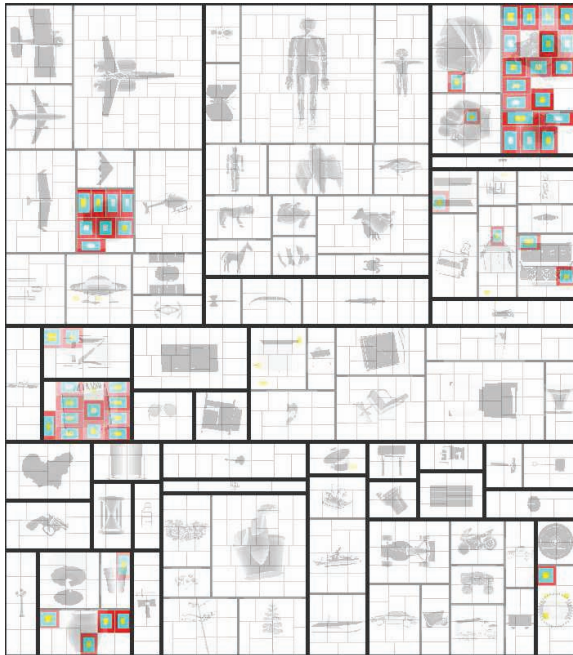


Figure 3: Same results as in Fig. 2 but with less features displayed. In particular, assuming we are primarily interested in items with similar geometric features and similar volume, the ratio of the volume of the item to the corresponding convex hull and the surface of objects are excluded from the visual representation. Now, the user can instantly identify the items that also have a similar volume to the query item as well (these are the items with a yellow component).

both in terms of the geometric features and the volume can easily be separated from items that have similar geometric features but differ in volume to the query item, as they are displayed with higher intensity color values in the channel that represents the relevance according to the volume.

5. User evaluation

In order to evaluate the developed multicriterial interface, user tests were executed, comparing it to a simple interface in which retrieved items are displayed in a serial manner. The already developed tool was used to present a serial list of results to the user, with the treemap visualization and the main controls hidden and navigation buttons added.

A group of 20 users was asked to perform 30 searches each on a set of 30 fixed items and explore the results. The experiment was repeated for both the serial interface and the treemap based interface. The users were given 5 minutes to familiarize themselves with the tools before executing the trials. The precision / recall scores computed based on the examined results for the two interfaces are presented in Fig. 4.




Figure 4: Precision / recall results for the serial and treemap based multifeature interface. The superiority of the treemap based interface is evident.

It is clear that the treemap based interface has better performance over the serial interface. In particular, precision is on average 13.1% higher for the same recall values for the treemap based interface compared to the serial interface. This should probably be attributed to the multicriterial presentation of results and the compact representation of results by the treemap.

6. Conclusion

A treemap based interface for visualization of search results according to multiple criteria has been presented. A merit of this method is that the relevance of items according to different search parameters can instantly be visually observed. Moreover, important groups of items can be instantly determined due to the hierarchical structure of the treemap. A critical factor for achieving this, has been the choice of the concentric rectangles layout as, compared to alternative arrangements, it has improved the clarity of the visual representation and maintains the visual cohesiveness of the features of each item.

Nevertheless, this arrangement is appropriate for a few features / criteria only, due to visual clutter. Appropriate choice of color or arrangement of features within the representation of a single item may be necessary to make this approach suitable for problems with more criteria. A possible extension could be the use of a 3D treemap [CHW09], with the third dimension being used to display the relevance according to the set of criteria at hand. In addition, the implemented system is appropriate only for small databases where all items can be simultaneously presented. In future extensions, we intend to implement hierarchical navigation through results with aggregated relevance scores being displayed at each level in order to guide the user in searching

larger databases. Finally, we intend to improve the search and ranking mechanisms by using relevance feedback methods [RL03]. These methods will be supported by interactive visual interfaces that will further assist the user in refining the search and allow the system to infer the subtleties of the provided search criteria.

Acknowledgements

This work was supported by the EU funded I-SEARCH IST STREP (FP7-248296).

References

- [BHvW99] BRULS M., HUIZING K., VAN WIJK J.: Squarified treemaps. In *In Proceedings of the Joint Eurographics and IEEE TCVG Symposium on Visualization* (1999), Press, pp. 33–42. 3
- [BKS*06] BUSTOS B., KEIM D., SAUPE D., SCHRECK T., VRANIC D.: An experimental effectiveness comparison of methods for 3d similarity search. *International Journal on Digital Libraries, Special 6*, 1 (2006), 39–54. 1
- [BKSS07] BUSTOS B., KEIM D., SAUPE D., SCHRECK T.: Content-based 3d object retrieval. *Computer Graphics and Applications, IEEE 27*, 4 (July-Aug. 2007), 22–27. 1
- [CDF09] CLARKSON E., DESAI K., FOLEY J.: Resultmaps: Visualization for search interfaces. *Visualization and Computer Graphics, IEEE Transactions on 15*, 6 (Nov.-Dec. 2009), 1057–1064. 1
- [CFR06] COPPERSMITH D., FLEISCHER L., RUDRA A.: Ordering by weighted number of wins gives a good ranking for weighted tournaments. In *SODA '06: Proceedings of the seventeenth annual ACM-SIAM symposium on Discrete algorithm* (New York, NY, USA, 2006), ACM, pp. 776–782. 2
- [CHW09] CHAUDHURI A., HAN-WEI S.: A self-adaptive treemap-based technique for visualizing hierarchical data in 3d. pp. 105–112. 5
- [FGL*08] FAN J., GAO Y., LUO H., KEIM D. A., LI Z.: A novel approach to enable semantic and visual image summarization for exploratory image search. In *MIR '08: Proceeding of the 1st ACM international conference on Multimedia information retrieval* (New York, NY, USA, 2008), ACM, pp. 358–365. 1
- [FKG*09] FAN J., KEIM D. A., GAO Y., LUO H., LI Z.: Justclick: personalized image recommendation via exploratory search from large-scale flickr images. *IEEE Transactions on Circuits and Systems for Video Technology 19*, 2 (2009), 273–288. 1
- [GPL*09] GAO Y., PENG J., LUO H., KEIM D., FAN J.: An interactive approach for filtering out junk images from keyword-based google search results. *IEEE Transactions on Circuits and Systems for Video Technology 19*, 12 (December 2009), 1851–1865. 1
- [JS91] JOHNSON B., SHNEIDERMAN B.: Tree-maps: a space-filling approach to the visualization of hierarchical information structures. In *VIS '91: Proceedings of the 2nd conference on Visualization '91* (Los Alamitos, CA, USA, 1991), IEEE Computer Society Press, pp. 284–291. 3
- [JWR00] JONES K. S., WALKER S., ROBERTSON S. E.: A probabilistic model of information retrieval: development and comparative experiments - part 1. *Inf. Process. Manage. 36*, 6 (2000), 779–808. 2
- [MAD*06] MADEMLIS A., AXENOPOULOS A., DARAS P., TZOVARAS D., STRINTZIS M. G.: 3d content-based search based on 3d krawtchouk moments. In *3DPVT '06: Proceedings of the Third International Symposium on 3D Data Processing, Visualization, and Transmission (3DPVT'06)* (Washington, DC, USA, 2006), IEEE Computer Society, pp. 743–749. 2
- [MDTS09] MADEMLIS A., DARAS P., TZOVARAS D., STRINTZIS M. G.: 3d object retrieval using the 3d shape impact descriptor. *Pattern Recognition 42*, 11 (2009), 2447–2459. 2
- [RL03] RUTHVEN I., LALMAS M.: A survey on the use of relevance feedback for information access systems. *Knowl. Eng. Rev. 18*, 2 (2003), 95–145. 6
- [SvZ09] SCHALEKAMP F., VAN ZUYLEN A.: Rank aggregation: Together we're strong. In *ALENEX* (2009), Finocchi I., Hersberger J., (Eds.), SIAM, pp. 38–51. 2
- [YPO03] YAP P., PARAMESRAN R., ONG S.: Image analysis by krawtchouk moments. 1367–1377. 2
- [ZDA*07] ZARPALAS D., DARAS P., AXENOPOULOS A., TZOVARAS D., STRINTZIS M. G.: 3d model search and retrieval using the spherical trace transform. *EURASIP J. Appl. Signal Process. 2007*, 1 (2007), 207–207. 2