Eurographics Workshop on 3D Object Retrieval (2016)
A. Ferreira, A. Giachetti, and D. Giorgi (Editors)

# 3D Objects Exploration: Guidelines for Future Research

S. Biasotti[1] and B. Falcidieno[1] and D. Giorgi[2] and M. Spagnuolo[1]

[1]CNR - IMATI, Consiglio Nazionale delle Ricerche, Italy
[2]CNR - ISTI, Consiglio Nazionale delle Ricerche, Italy

**Abstract**

*Search engines provide the interface to interact with 3D object repositories, which are rapidly growing in both number and size. This position paper presents the current state of the art on 3D dataset navigation and 3D model retrieval. We discuss a number of challenges we consider as the main points to be tackled for developing effective 3D object exploration systems.*

Categories and Subject Descriptors (according to ACM CCS): Computer Graphics [I.3.6]: Methodology and Techniques—

## 1. Introduction

The use of search engines is now the default way in which most people engage with information, as the vast majority of information is available today in digital form. A new wave of information has been underway for a few years: low-cost 3D scanners and sensors, 3D printers, and powerful modelling software have made 3D models familiar to an increasing audience [SF09]. 3D digital representations of objects have been penetrating applied domains spanning from entertainment to life sciences, cultural heritage, industrial design, manufacturing and urban planning [CMSF11]. Repositories of 3D objects are rapidly growing in both number and size [BLMS14, Spa16]. They are a valuable resource for both consolidated and new uses, as it is easier to find something, and possibly modify it, than to create it from scratch [TGY*09]. Therefore, exploring 3D repositories and retrieving appropriate models has been drawing attention from researchers for over fifteen years.

### 1.1. A brief historical excursus on 3D retrieval: from geometric descriptors to the semantics of 3D collections

As traditional information access techniques were shown inadequate to deal with 3D content, the first *content-based* approaches were presented at the beginning of the 2000s [VH01]. They focused on the definition of *shape descriptors* as signatures synthesizing the geometric content of 3D models [BKS*05, TV08], and of *similarity measures* for matching descriptors [SJ99, BCBB15]. The result was the introduction of the first 3D search engine prototypes [SMKF04]. Over the years, descriptors and similarity measures were refined to deal with geometric invariances, from rigid transformations to isometries, according to the application at hand [BBK08, BFGS14]. Appearance was also taken into account beside geometry, resulting in techniques for the retrieval of *textured* 3D models, which are common for example in cultural heritage applications [BCGS13, BCFS15].

Afterward, researchers realized that there was much more to the problem of 3D retrieval than just geometry and appearance. Gradually, the focus shifted from the concept of geometric *shape* to that of *semantics*, as the meaning or *functionality* of a 3D object in a given context. Many descriptors were proposed to infer semantic (high-level) information from low-level properties [LMS13]. Moreover, users can be interested in exploring different aspects of 3D objects, thus calling for modular descriptors which could be fine-tuned to the problem at hand [BDF*08]. Also, as similarity is a cognitive process, the user's intent during the search should be included in the loop [SWS*00]: methods for relevance feedback were introduced for 3D object search [LMT05, GFSF10] and the first 3D search engines supporting user interaction appeared quite recently [Hea11].

The most recent path of research aims to go beyond the retrieval of individual objects in a 3D collection, by supporting a *collection-level* 3D content analysis: namely, the *navigation* across an entire collection, the abstraction of family of shapes and the summarization of shape variability [LHLF15], up to 3D collection exploration for content-based modelling and editing [ZLDM15].

### 1.2. The future of 3D data exploration

Despite the considerable efforts, 3D retrieval and 3D dataset navigation remain challenging problems. The issues still unsolved, or not receiving enough attention, include the definition of the concepts of similarity itself, the building of user-friendly interfaces for content-based 3D object search, the development of effective visualization techniques for 3D search results, and the handling of real-time response of techniques and algorithms.

In this position paper, we identify some of the main challenges for searching 3D collections, and suggest seven research guidelines for designing the 3D navigation systems of the future.



DOI: 10.2312/3dor.20161081

## 2. The road ahead

### 2.1. Challenge one: Supporting search intent

In Information Retrieval, the strategies for navigating an electronic collection can be broadly classified into two categories: *analytical* and *browsing* [KRD*15, MS88]. Analytical strategies serve to retrieve specific facts, whereas browsing supports exploratory information-seeking tasks. In 3D search and retrieval, from the very beginning the research focused on analytical strategies for matching user queries to 3D database objects. The aim was to support immediate navigation to desired objects, assuming that users had a clear idea of their target from the start. Nevertheless, users often navigate towards information targets via iterative steps using cues from the environment, and also change their perspective on the search task as they interact with the content of the environment (see [KRD*15], even if applied to a different context). Also, as observed by Bates [Bat89], individuals pick up bits and pieces of information as they navigate through an information space. Finally, browsing strategies help circumventing difficulties in querying the system (the well-known *vocabulary problem* or *page-zero problem*).

Therefore, future systems should be able both to answer specific queries and to support continuous, iterative navigation of the 3D resources. Browsing strategies are supported by the recent proposals in [KFLCO13], which arranges the objects in a 3D dataset in a dynamic map that the user can spatially navigate; and [HSS*13], which arranges the objects as either a hierarchical tree the user can traverse starting from any object in a leaf node, or as a set of concentric circles around a specific query.

### 2.2. Challenge two: Supporting faceted search

3D shapes can vary in several ways, and users may be interested in exploring different types of variations; even within a single exploration session, a user may want to define the exploration space using multiple different attributes. *Facets* are quantitative or categorical aspects which are of importance to a user [JGZ*11]. *Faceted browsing* allows users to specify filters to find subsets of items sharing specific desired characteristics. In the context of 3D objects search, users may be interested in different geometric or semantic attributes.

Therefore, a search system should allow users to dynamically prescribe multiple *facets* to drive the exploration. This requirement impacts on both the building of interfaces and the definition of 3D descriptors. On the one hand, faceted browsing interfaces have to be defined to support users specify their filters, as in [KLM*12], where the filters simply are regions of interest on the surface (e.g., the back of a chair). Also, as users may not know a priori which properties they are interested in, the interface should help them decide on which filters to specify. On the other hand, the search system should be endowed with 3D descriptors able to encode different shape and semantic properties. Recent mathematical theories such as *persistent topology* can help to define modular descriptors which adapt to different shape properties: they offer a common framework in which different functions can be plugged in to describe different properties of an object [BDF*08].

### 2.3. Challenge three: Offering contextual cues to support search refinement

Once users have started navigating, a system should support query refinement. Presenting the users with the context in which the retrieved content is encoded may help them refining their search. For example, *Degree-of-Interest* (DOI) visualization techniques [Fur86] highlight items of interest along with subsets of objects which provide explanatory context. This idea is explored in [HSS*13], which uses *quartets* of objects to organize a heterogeneous collection of shapes and allow users to gain an overview of it. Inspiring ideas can be found in [BTP*12], which suggests a focus-plus-context principle for the dynamic exploration of a (already) clustered dataset. A selected thumbnail image determines the layout (size and position) of the other images, which are close and big if semantically similar to the focus image, while small and stacked in far away piles if dissimilar.

Another possibility is to look at the literature on serendipitous finding. *Serendipity* refers to the ability of connecting unexpectedly found information with what you know and what you are possibly looking for. [DCW11] suggests to encourage serendipity by showing users items which "share unusual facets or relate to one's previous interactions". In this direction, an interesting proposal is the *fit&diverse* paradigm in [XZCOC12], which presents users with both objects which fit their target and diverse objects to create surprising and inspiring suggestions.

### 2.4. Challenge four: Relevance feedback and user engagement

3D shape search is an interaction process between the user, the 3D content and its semantics. Smeulders et al. [SWS*00] brought the *semantic gap* into focus, namely the gap between the visual data information and the meaning of the data for the user. Indeed, as observed by Koenderink [Koe90], things possess a shape for the observer: in the observer's mind, the perception of a phenomenon joins existing concepts and enables recognition and similarity assessment. Similarity is a cognitive process, which depends upon the observer. Therefore, it is mandatory to make the user active in the search process [DJLW08]. To "include a human in the loop", *relevance feedback* techniques enable users to feed the system with their opinion so that it refines its answer accordingly. Relevance Feedback dates back to 1971, when it was introduced by Rocchio for text document retrieval [Roc71]. Then, in the 1990s Rui et al. brought it to the Computer Vision Community [RWOM98]. It was only recently that relevance feedback started drawing the attention of the 3D retrieval community [LMT05, PPT*08, GFSF10].

The 3D search system of the future should engage users directly, favour a continuous, active dialogue, and make it easy for users to provide feedback about suggested items. Also, the system should be able to analyse users' behaviour and infer their preferences, without posing a burden on them, that is, to infer the *ostensive relevance* of items directly from the user interaction with the content [Cam96]. Two examples in this direction are [KFLCO13], in which a dynamic map is progressively filled with the 3D objects in the dataset according to the user's navigation tendencies; and [HSS*13], in which the system maintains the exploration history of the collection in a separate window, allowing simple understanding of the exploration path and quick backtracking.

### 2.5. Challenge five: Scalability and real-time interaction

Whereas the first 3D datasets were made of hundreds or, in some cases, thousands of objects [SMKF04], the current repositories may contain up to millions of models, like the Trimble warehouse†. Also, the increased availability of 3D acquisition facilities, including crowd-based photogrammetric methods, or cheap sensors like Microsoft Kinect, large-scale shape retrieval will become important [SBS*15]. Therefore, the scalability of the exploration techniques with respect to the size of the dataset is mandatory to support both real-time result visualization and real-time user interaction. Traditional query modalities based on similarity matrices (quadratic in terms of the dataset elements) are not suitable when dealing with hundred thousands of models, like in [GCL*15]. Possible solutions are optimization via parallel implementation of some steps [KLM*12]; two-stage strategies with off-line dataset clustering preceding on-line dataset exploration [HSS*13]; reducing the number of on-line computations by taking into account only the neighbourhood of the current objects of interest [KFLCO13].

Finally, we observe that scalable methods should be also robust against degenerate models, which can represent a large fraction of large-scale 3D repositories [SBS*15].

### 2.6. Challenge six: Supporting navigation across non homogeneous datasets

Future 3D search systems should enable exploration across different data, such as 3D objects, images, videos, as well as textual metadata. For instance, in the Cultural Heritage domain, collections of 3D models are equipped with multiple information like images representing details of relevance, archaeological descriptions and metadata on the provenance, the age, the style, etc. Therefore, it would be important to uncouple the representation format of a resource from the query modality by which to find it (for instance, accessing a 3D model by querying the system with 2D images), and to offer an integrated approach to the management of 2D, 3D data and semantic metadata.

Under the assumption that 3D models and 2D images provide complementary information about the same objects or object classes [HOM15], recent approaches in both computer graphics and computer vision exploit cross-domain information to improve the performance of techniques in each domain, including part-based model synthesis [XZCO*11], consistent 3D shape segmentation [WAvK*12] and co-alignment, object detection and camera pose estimation [AME*14], object shape estimation in images and learning efficient classifiers on 3D shapes for better text-based search [HOM15]. Similar ideas could be implemented to support representation-independent browsing techniques of hybrid datasets.

### 2.7. Challenge seven: Domain-specific 3D search

While the entertainment industry has been the driving force behind the explosion of applications and technological advances in 3D media, 3D models no longer are just fancy-looking objects: 3D repositories are essential knowledge carriers and may represent a huge economic factor in many applied and scientific domains, which include design and manufacturing, serious gaming and simulation, cultural heritage and archaeology, medicine and bioinformatics, geographic information systems and environment [CMSF11]. Each domain has its own needs and peculiarities: searching CAD objects in a repository is a different problem than navigating a catalogue of ancient Greek vessels. In addition, in the same domain an object can be seen from different points of view, which strongly affect the concept of similarity (e.g. in product modelling: design, manufacturing, packaging, etc.).

Therefore, future 3D search engines should adapt their answers to both the application domain and the user's point of view. Again, this poses requirements on both the choice of descriptors and metrics and the interface design. The choice of descriptors and metrics should reflect the characteristics of the field, for example supporting partial matching for the search of fragments in cultural heritage datasets [SPS15]. The interface should offer facilities, which enhance the user domain environment, experience and professional skills, for example supporting medical doctors with the integration of their medical knowledge with the content-based search tools offering a unified user's interface [BCPS15].

† https://3dwarehouse.sketchup.com/

## 3. Conclusions

We have presented seven research challenges which we hope may help inspire further research and serve as guidelines for designing effective and interactive 3D search systems. Of course, our list is not exhaustive. New challenges are coming out which lie at the boundary between disciplines. Two examples are data mining for the integration of massive quantities of user behaviours to improve search algorithms, and tools for social and collaborative search. We believe solving these problems will require the cross-fertilization between different fields such as Computer Vision, Computer Graphics, Machine Learning, Cognitive Science and Human-Computer Interaction.

## References

[AME*14] AUBTY M., MATURANA D., EFROS A. A., RUSSELL B., SIVIC J.: Seeing 3d chairs: exemplar part-based 2d-3d alignment using a large dataset of cad models. In *IEEE CVPR Proc.* (2014). 3

[Bat89] BATES M. J.: The design of browsing and berrypicking techniques for the online search interface. *Online Review 13* (1989), 407–424. 2

[BBK08] BRONSTEIN A., BRONSTEIN M., KIMMEL R.: *Numerical Geometry of Non-Rigid Shapes*, 1 ed. Springer P. C., Inc., 2008. 1

[BCBB15] BIASOTTI S., CERRI A., BRONSTEIN A., BRONSTEIN M.: Recent trends, applications, and perspectives in 3d shape similarity assessment. *Computer Graphics Forum* (2015). 1

[BCFS15] BIASOTTI S., CERRI A., FALCIDIENO B., SPAGNUOLO M.: 3D artifacts similarity based on the concurrent evaluation of heterogeneous properties. *Journal on Computing and Cultural Heritage (JOCCH) 8*, 4 (2015), 19. 1

[BCGS13] BIASOTTI S., CERRI A., GIORGI D., SPAGNUOLO M.: PHOG: Photometric and Geometric Functions for Textured Shape Retrieval. *Computer Graphics Forum 32*, 5 (2013), 13–22. 1

[BCPS15] BANERJEE I., CATALANO C. E., PATANÉ G., SPAGNUOLO M.: Semantic annotation of 3D anatomical models to support diagnosis and follow-up analysis of musculoskeletal pathologies. *International journal of computer assisted radiology and surgery* (2015), 1–14. 3

[BDF*08] 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 (2008), 1–87. 1, 2

[BFGS14] Biasotti S., Falcidieno B., Giorgi D., Spagnuolo M.: Mathematical tools for shape analysis and description. *Synthesis Lectures on Computer Graphics and Animation 6*, 2 (2014), 1–138. 1

[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. 1

[BLMS14] Biasotti S., Laga H., Mortara M., Spagnuolo M.: Reasoning about shape in complex datasets: Geometry, structure and semantics. In *Tutorial Notes* (2014), The Eurographics Association. 1

[BTP*12] Brivio P., Tarini M., Ponchio F., Cignoni P., Scopigno R.: Pilebars: Scalable dynamic thumbnail bars. In *VAST 2012 Symp. Proc.* (2012), Eurographics Association, pp. 49–56. 2

[Cam96] Campbell I.: *The ostensive model of developing information needs*. PhD thesis, University of Glasgow, 1996. 2

[CMSF11] Catalano C. E., Mortara M., Spagnuolo M., Falcidieno B.: Semantics and 3d media: Current issues and perspectives. *Computers & Graphics 35*, 4 (2011), 869–877. 1, 3

[DCW11] Dork M., Carpendale S., Williamson C.: The information flaneur : A fresh look at information seeking. In *ACM CHI Proc.* (2011), p. 1215Ű1224. 2

[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). 2

[Fur86] Furnas G. W.: Generalized fisheye views. *ACM SIGCHI Bulletin 17*, 4 (1986), 16–23. 2

[GCL*15] Gao L., Cao Y.-P., Lai Y.-K., Huang H.-Z., Kobbelt L., Hu S.-M.: Active exploration of large 3D model repositories. *IEEE Transactions on Visualization and Computer Graphics 99*, PrePrints (2015), 1. 3

[GFSF10] Giorgi D., Frosini P., Spagnuolo M., Falcidieno B.: 3d relevance feedback via multilevel relevance judgements. *The Visual Computer 26*, 10 (2010), 1321–1338. 1, 2

[Hea11] Hearst M. A.: 'natural' search user interfaces. *Commun. ACM 54*, 11 (Nov. 2011), 60–67. 1

[HOM15] Hueting M., Ovsjanikov M., Mitra N.: Crosslink: Joint understanding of image and 3D model collections through shape and camera pose variations. *ACM SIGGRAPH Asia 2015* (2015). 3

[HSS*13] Huang S.-S., Shamir A., Shen C.-H., Zhang H., Sheffer A., Hu S.-M., Cohen-Or D.: Qualitative organization of collections of shapes via quartet analysis. *ACM Trans. Graph. 32*, 4 (July 2013), 71:1–71:10. 2, 3

[JGZ*11] Jetter H.-C., Gerken J., Zöllner M., Reiterer H., Milic-Frayling N.: Materializing the query with facet-streams: A hybrid surface for collaborative search on tabletops. In *CHI 2011* (May 2011), ACM. 2

[KFLCO13] Kleiman Y., Fish N., Lanir J., Cohen-Or D.: Dynamic maps for exploring and browsing shapes. *Computer Graphics Forum 32*, 5 (2013), 187–196. 2, 3

[KLM*12] Kim V. G., Li W., Mitra N. J., DiVerdi S., Funkhouser T.: Exploring collections of 3D models using fuzzy correspondences. *ACM T. Graphic. 31*, 4 (2012), 54:1–54:11. 2, 3

[Koe90] Koenderink J. J.: *Solid shape*. MIT Press, Cambridge, MA, USA, 1990. 2

[KRD*15] Kairam S., Riche N. H., Drucker S., Fernandez R., Heer J.: Refinery: Visual exploration of large, heterogeneous networks through associative browsing. *Computer Graphics Forum 34*, 3 (2015), 301–310. 2

[LHLF15] Liu T., Hertzmann A., Li W., Funkhouser T.: Style compatibility for 3D furniture models. *ACM Transactions on Graphics (Proc. SIGGRAPH) 34*, 4 (Aug. 2015). 1

[LMS13] Laga H., Mortara M., Spagnuolo M.: Geometry and context for semantic correspondences and functionality recognition in man-made 3d shapes. *ACM Trans. Graph. 32*, 5 (Oct. 2013), 150:1–150:16. 1

[LMT05] Leifman G., Meir R., Tal A.: Semantic-oriented 3d shape retrieval using relevance feedback. *The Visual Computer 21*, 8-10 (2005), 865–875. 1, 2

[MS88] Marchionni G., Shneiderman B.: Finding facts vs. browsing knowledge in hypertext systems. *Computer 21*, 1 (1988), 70–80. 2

[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. 2

[Roc71] Rocchio J.: *Relevance feedback in Information Retrieval*. Prentice Hall, 1971. 2

[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. 2

[SBS*15] Sipiran I., Bustos B., Schreck T., Bronstein A. M., Bronstein M., Castellani U., Choi S., Lai L., Li H., Litman R., Sun L.: Scalability of non-rigid 3d shape retrieval. In *3DOR* (2015). 3

[SF09] Spagnuolo M., Falcidieno B.: 3D media and the semantic web. *IEEE Intelligent Systems 24*, 2 (2009), 90–96. 1

[SJ99] Santini S., Jain R.: Similarity measures. *IEEE Trans. Pattern Anal. Mach. Intell. 21*, 9 (Sept. 1999), 871–883. 1

[SMKF04] Shilane P., Min P., Kazhdan M., Funkhouser T.: The Princeton Shape Benchmark. In *Shape Modeling International* (2004), IEEE Computer Society, pp. 167–178. 1, 3

[Spa16] Spagnuolo M.: Shape 4.0: 3D shape modeling and processing using semantics. *Computer Graphics and Applications, IEEE 36*, 1 (2016), 92–96. 1

[SPS15] Savelonas M. A., Pratikakis I., Sfikas K.: Partial 3D Object Retrieval combining Local Shape Descriptors with Global Fisher Vectors. In *3DOR 2015* (2015). 3

[SWS*00] Smeulders A. W. M., Worring M., Santini S., Gupta A., Jain R.: Content-based image retrieval at the end of the early years. *IEEE Trans. Pattern Anal. Mach. Intell. 22*, 12 (Dec. 2000), 1349–1380. 1, 2

[TGY*09] Talton J. O., Gibson D., Yang L., Hanrahan P., Koltun V.: Exploratory modeling with collaborative design spaces. *ACM Trans. Graph. 28*, 5 (Dec. 2009), 167:1–167:10. 1

[TV08] Tangelder J. W. H., Veltkamp R. C.: A survey of content based 3D shape retrieval methods. *Multimedia Tools Appl. 39*, 3 (2008), 441–471. 1

[VH01] Veltkamp R. C., Hagendoorn M.: State-of-the-Art in shape matching. In *Principles of Visual Information Retrieval*, Lew M., (Ed.). Springer-Verlag, 2001, pp. 87–119. 1

[WAvK*12] Wang Y., Asafi S., van Kaick O., Zhang H., Cohen-Or D., Chen B.: Active co-analysis of a set of shapes. *ACM Trans. Graph. 31*, 6 (Nov. 2012), 165:1–165:10. 3

[XZCO*11] Xu K., Zheng H., Cohen-Or D., Liu L., Xiong Y.: Photo-inspired model-driven 3D object modeling. *ACM SIGGRAPH 2011* (2011), 80:1–80:10. 3

[XZCOC12] Xu K., Zhang H., Cohen-Or D., Chen B.: Fit and diverse: Set evolution for inspiring 3d shape galleries. *ACM Trans. Graph. 31*, 4 (July 2012), 57:1–57:10. 2

[ZLDM15] Zheng Y., Liu H., Dorsey J., Mitra N. J.: Ergonomics-inspired reshaping and exploration of collections of models. *IEEE Transactions on Visualization and Computer Graphics* (2015). 1