ShapeBench: A new approach to benchmarking local 3D shape descriptors

Bart Iver van Blokland

Abstract

"The ShapeBench evaluation methodology is proposed as an extension to the popular Area Under Precision-Recall Curve (PRC/AUC) for measuring the matching performance of local 3D shape descriptors. [...]"

Overview

- Relevant theory
- The ShapeBench benchmark
- Results

Relevant theory

Local 3D shape descriptors - A 2D analogy

(a)

David G. Lowe, “Distinctive Image Features from Scale-Invariant Keypoints,” International Journal of Computer Vision 60, no. 2 (November 1, 2004): 91–110

Local 3D shape descriptors - A 2D analogy

- Perform keypoint localization

(a)

(d)

David G. Lowe, “Distinctive Image Features from Scale-Invariant Keypoints,” International Journal of Computer Vision 60, no. 2 (November 1, 2004): 91–110

Local 3D shape descriptors - A 2D analogy

- Perform keypoint localization

(a)

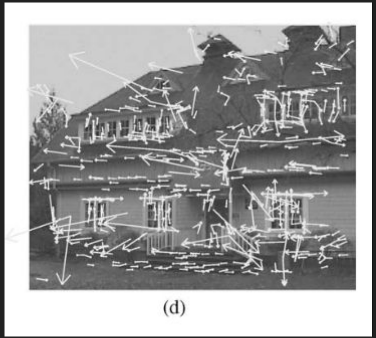

(d)

- Compute local image descriptor using area around keypoint

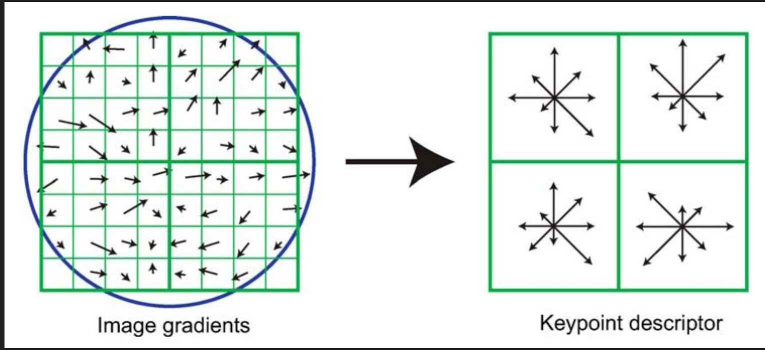

David G. Lowe, “Distinctive Image Features from Scale-Invariant Keypoints,” International Journal of Computer Vision 60, no. 2 (November 1, 2004): 91–110

Local 3D shape descriptors

The support radius determines the area to consider around the key point.

Is defined on a per-method basis in this experiment

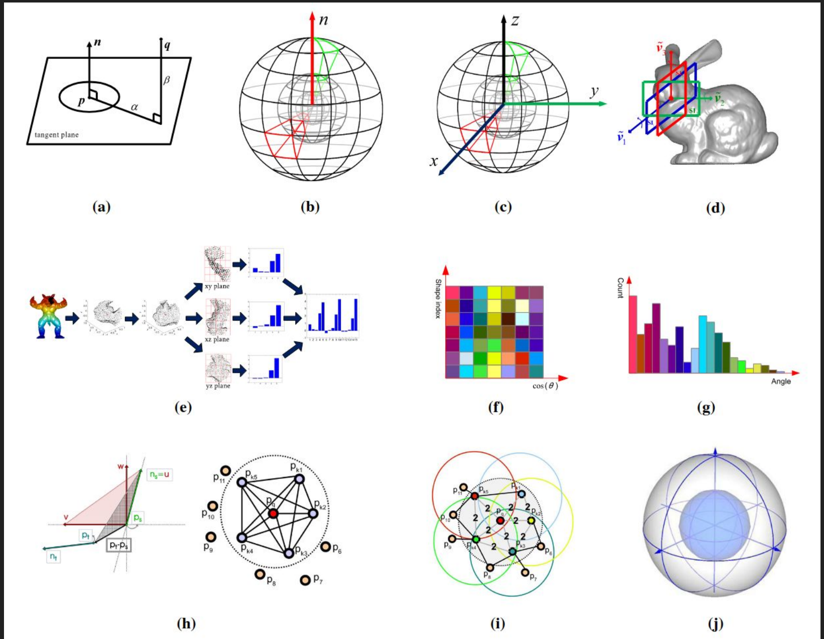

Yulan Guo et al., “A Comprehensive Performance Evaluation of 3D Local Feature Descriptors,” International Journal of Computer Vision 116, no. 1 (January 1, 2016), 69, fig. 1

Local 3D shape descriptors

The support radius determines the area to consider around the key point.

Is defined on a per-method basis in this experiment

Example of a descriptor:

Spin Image - Cylindrical coordinates of nearby points with respect to the key point and its normal vector

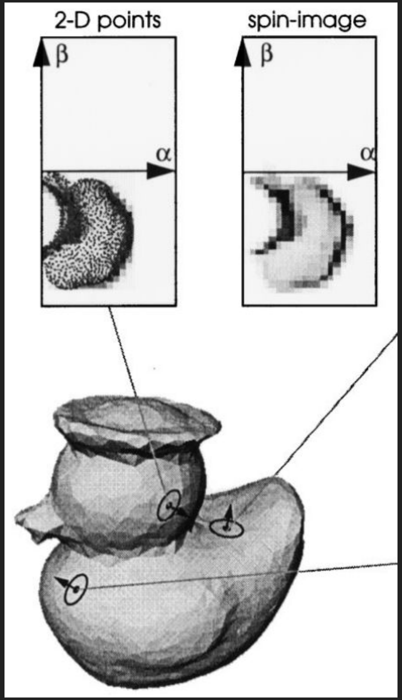

A. E. Johnson and M. Hebert, “Surface Matching for Object Recognition in Complex Three-Dimensional Scenes,” Image and Vision Computing 16, no. 9 (July 1, 1998), 638, fig. 3

Comparing descriptors: Nearest neighbor distance ratio (NNDR)

- For a descriptor, find its closest and next-closest neighbor in the reference set. (Closest as in shortest Euclidean distance)

Comparing descriptors:
Nearest neighbor distance ratio (NNDR)

- For a descriptor, find its closest and next-closest neighbor in the reference set. (Closest as in shortest Euclidean distance)
- The NNDR is the ratio between the distance to these
- Set a threshold τ for deciding match vs. no match

Area Under Precision-Recall Curve (PRC/AUC)

- Compute local feature descriptors for key points in the models
- Use NNDR to consider matches in reference set
- Plot the Precision-Recall curve (PRC) for τ [0, 1]

Area Under Curve (AUC) is simply the area under the PRC

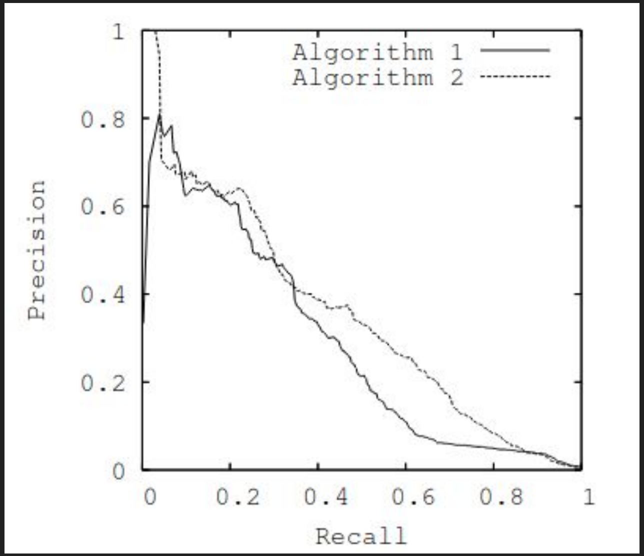

Jesse Davis and Mark Goadrich, “The Relationship between Precision-Recall and ROC Curves,” in Proceedings of the 23rd International Conference on Machine Learning, ICML ’06 (New York, NY, USA: Association for Computing Machinery, 2006), 234, fig. 1

The ShapeBench benchmark

ShapeBench

- Uses procedural generation to create a large dataset
- Introduces an extension of PRC/AUC, called the DDI

Dataset

Dataset Information				Used in Evaluation					
Dataset	**Model Set**	**Models**	**Scenes**	[10]	[12]	[9]	[13]	[14]	[11]
Bologna 3D Retrieval (B3R) [15]	Stanford	6	18	Yes	-	Yes	Yes	Yes	Yes
Random Views [15]	Stanford	6	36	-	-	Yes	-	-	-
Bologna Dataset 1&2 - Stanford [16]	Stanford	6	45	-	Yes	Yes	-	-	-
UWA 3D Modelling	UWA	4	75	Yes	-	-	*Yes*[1]	Yes	-
UWA Object Retrieval [17, 18]	UWA	5	50	Yes	Yes	-	Yes	Yes	Yes
Bologna Dataset 3 - SpaceTime Stereo [16]	Kinect (+ clutter)	8	15	-	-	Yes	-	-	Yes
Bologna Dataset 5 - Kinect [19]	Kinect (+ clutter)	6	16	Yes	-	*Yes*[1]	-	-	-
Bologna Object Recognition	Kinect (+ clutter)	6	17	-	Yes	-	-	-	-
Bologna Mesh Registration	Kinect	6	95	Yes	-	-	Yes	-	-
Queens LiDAR [20]	Queens	5	63	-	Yes	-	-	-	-
7-scenes [21]	7-scenes	7	n/a	-	-	Yes	-	-	-
DTU [22]	DTU	45	3,204	-	-	-	-	-	-
ShapeNetCore [23]	ShapeNetCore	51,300	n/a	-	-	-	-	-	-
ABC [24]	ABC	1,000,000	n/a	-	-	-	-	-	-
Objaverse [25]	Objaverse	798,759	n/a	-	-	-	-	-	-

Table 1: An overview over datasets used for the evaluations in a number of recent papers, as well as some examples of larger datasets. Datasets that were *not* used in a particular evaluation are marked with a hyphen (-) for visual clarity. All datasets with equivalent *model set* names use the same (sub)set of models.

Bart Iver van Blokland, "ShapeBench: A New Approach to Benchmarking Local 3D Shape Descriptors," Computers & Graphics 124 (November 1, 2024), 3, Table. 1

Procedural generation

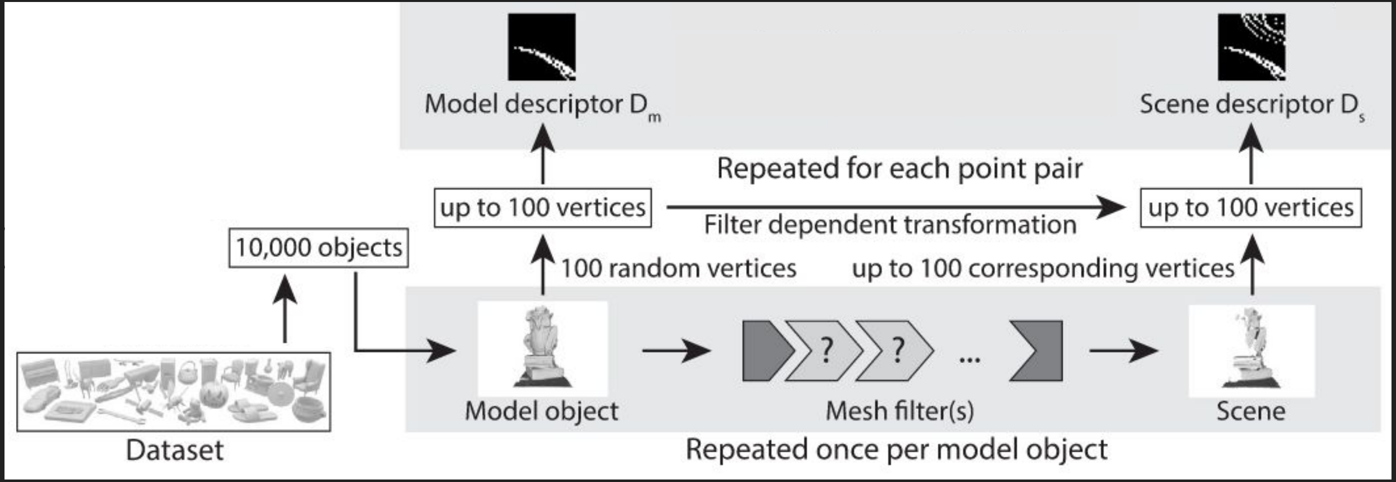

The Descriptor Distance Index (DDI)

PRC is affected by multiple valid matches.

To solve this, DDI uses a reference set R created from a large set of descriptors.

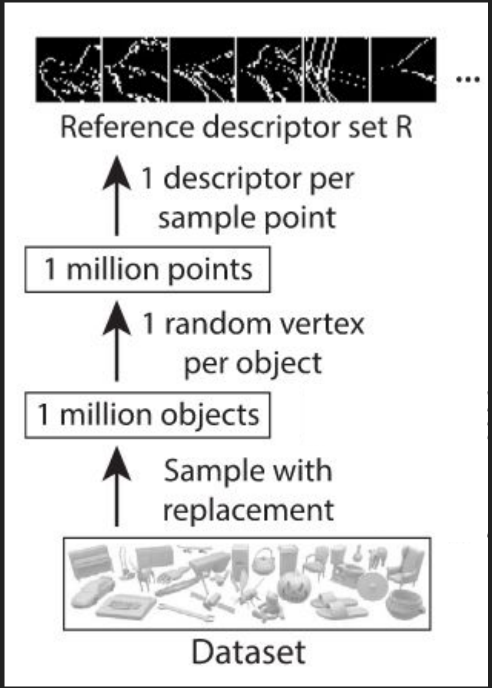

The Descriptor Distance Index (DDI)

PRC is affected by multiple valid matches.

To solve this, DDI uses a reference set R created from a large set of descriptors.

DDI(f_1, f_2) is the number of descriptors from R that are closer to f_1 than f_2 is.

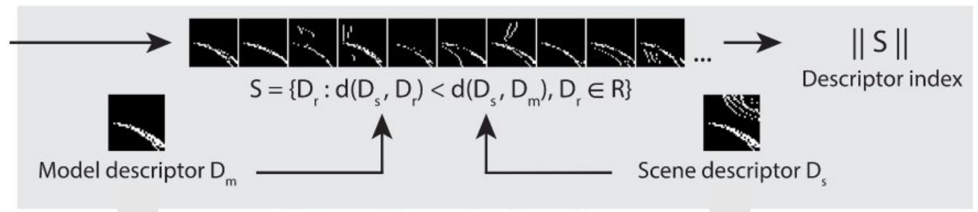

Filters

- Clutter
  - Piles random objects around the model

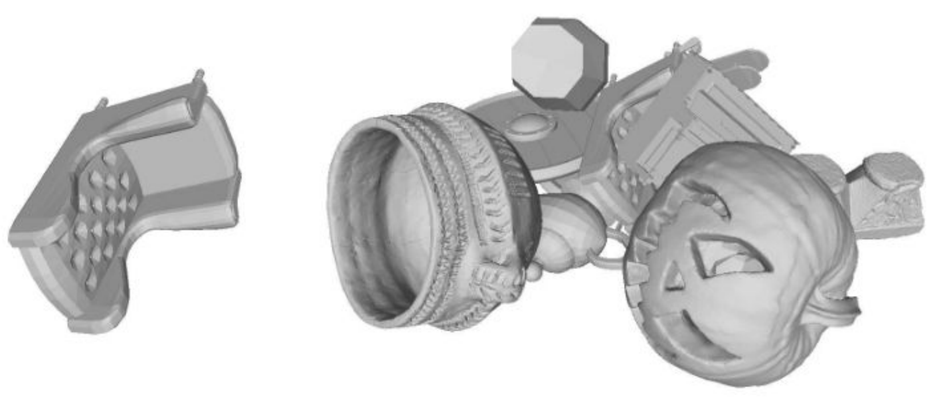

(a) Original (b) Filtered

Filters

- Clutter
- Occlusion
  - Removes vertices that are not visible from a random view direction

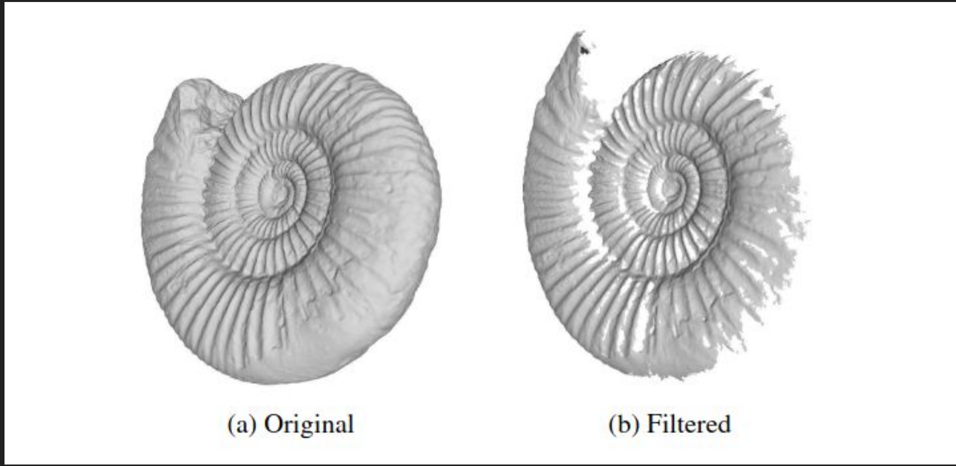

Bart Iver van Blokland, “ShapeBench: A New Approach to Benchmarking Local 3D Shape Descriptors,” Computers & Graphics 124 (November 1, 2024), 15, fig. 19

Filters

- Clutter
- Occlusion
- Alternate triangulation
  - Applies a mesh smoothing algorithm to displace vertices

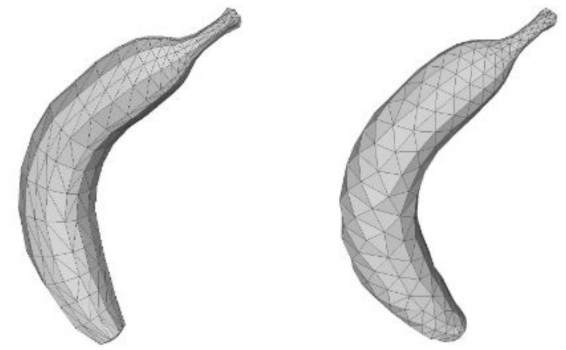

(a) Original (b) Filtered

Bart Iver van Blokland, “ShapeBench: A New Approach to Benchmarking Local 3D Shape Descriptors,” Computers & Graphics 124 (November 1, 2024), 15, fig. 20

Filters

- Clutter
- Occlusion
- Alternate triangulation
- Normal vector deviation
  - Randomly deviates all normals

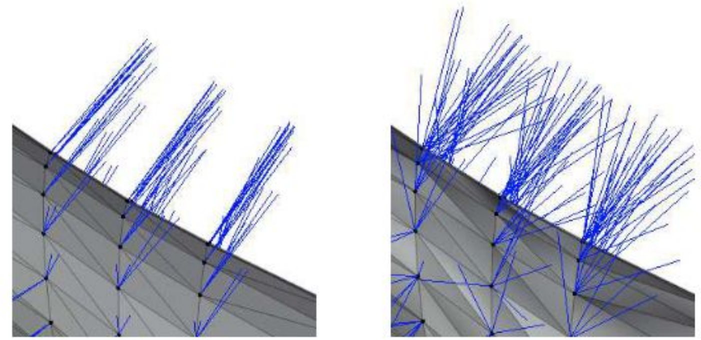
(a) Original (b) Filtered

Bart Iver van Blokland, “ShapeBench: A New Approach to Benchmarking Local 3D Shape Descriptors,” Computers & Graphics 124 (November 1, 2024), 15, fig. 21

Filters

- Clutter
- Occlusion
- Alternate triangulation
- Normal vector deviation
- Support radius deviation
    - Scales the model by a random factor to simulate a deviated support radius

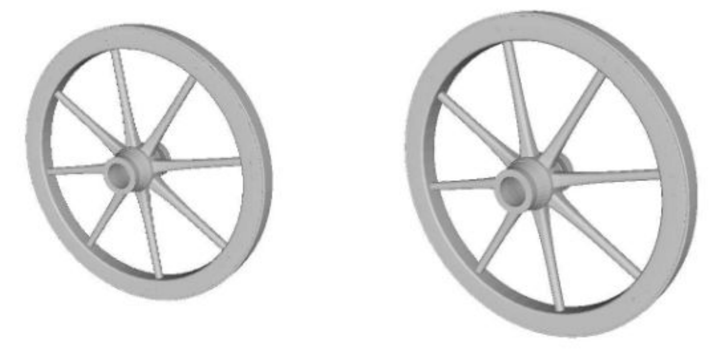

(a) Original (b) Filtered

Filters

- Clutter
- Occlusion
- Alternate triangulation
- Normal vector deviation
- Support radius deviation
- Gaussian noise
    - Simulates inaccurate scans by applying noise to vertex positions

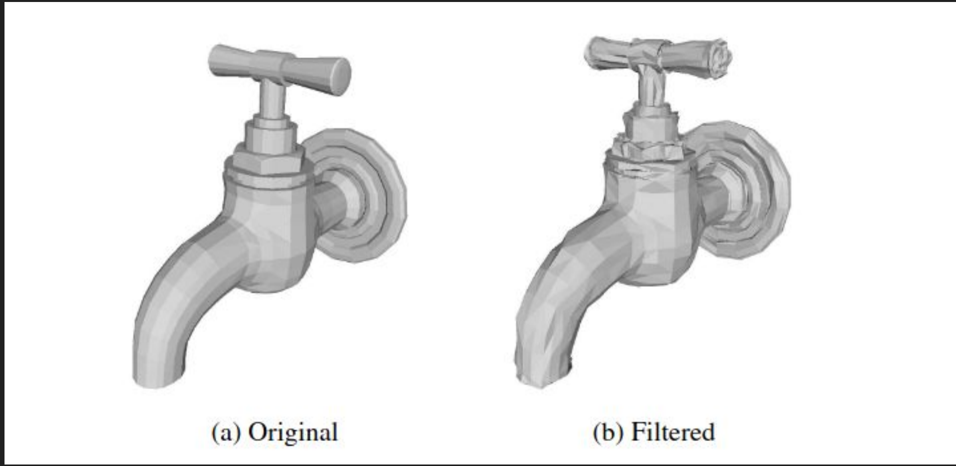

Bart Iver van Blokland, “ShapeBench: A New Approach to Benchmarking Local 3D Shape Descriptors,” Computers & Graphics 124 (November 1, 2024), 15, fig. 23

Filters

- Clutter
- Occlusion
- Alternate triangulation
- Normal vector deviation
- Support radius deviation
- Gaussian noise
- Alternate mesh resolution
    - Reconstructs the mesh from a low resolution render from a random view direction

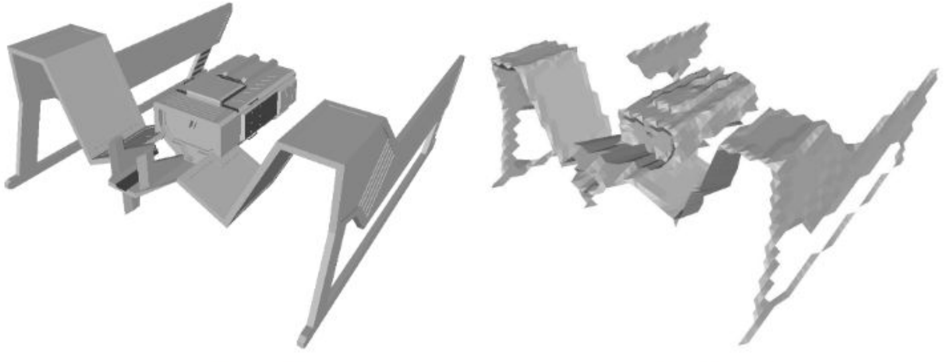

(a) Original (b) Filtered

Bart Iver van Blokland, “ShapeBench: A New Approach to Benchmarking Local 3D Shape Descriptors,” Computers & Graphics 124 (November 1, 2024), 15, fig. 24

Filters

- Clutter
- Occlusion
- Alternate triangulation
- Normal vector deviation
- Support radius deviation
- Gaussian noise
- Alternate mesh resolution
- Combination

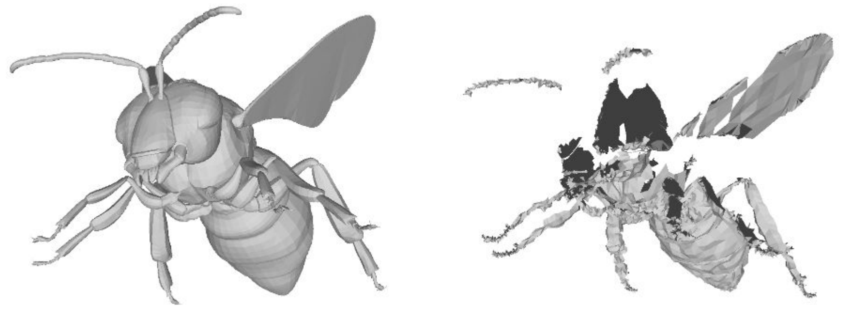

(a) Original (b) Filtered

Fig. 27: Example output of the occlusion filter followed by the Gaussian noise filter.

Bart Iver van Blokland, “ShapeBench: A New Approach to Benchmarking Local 3D Shape Descriptors,” Computers & Graphics 124 (November 1, 2024), 15, fig. 27

Benchmark overview

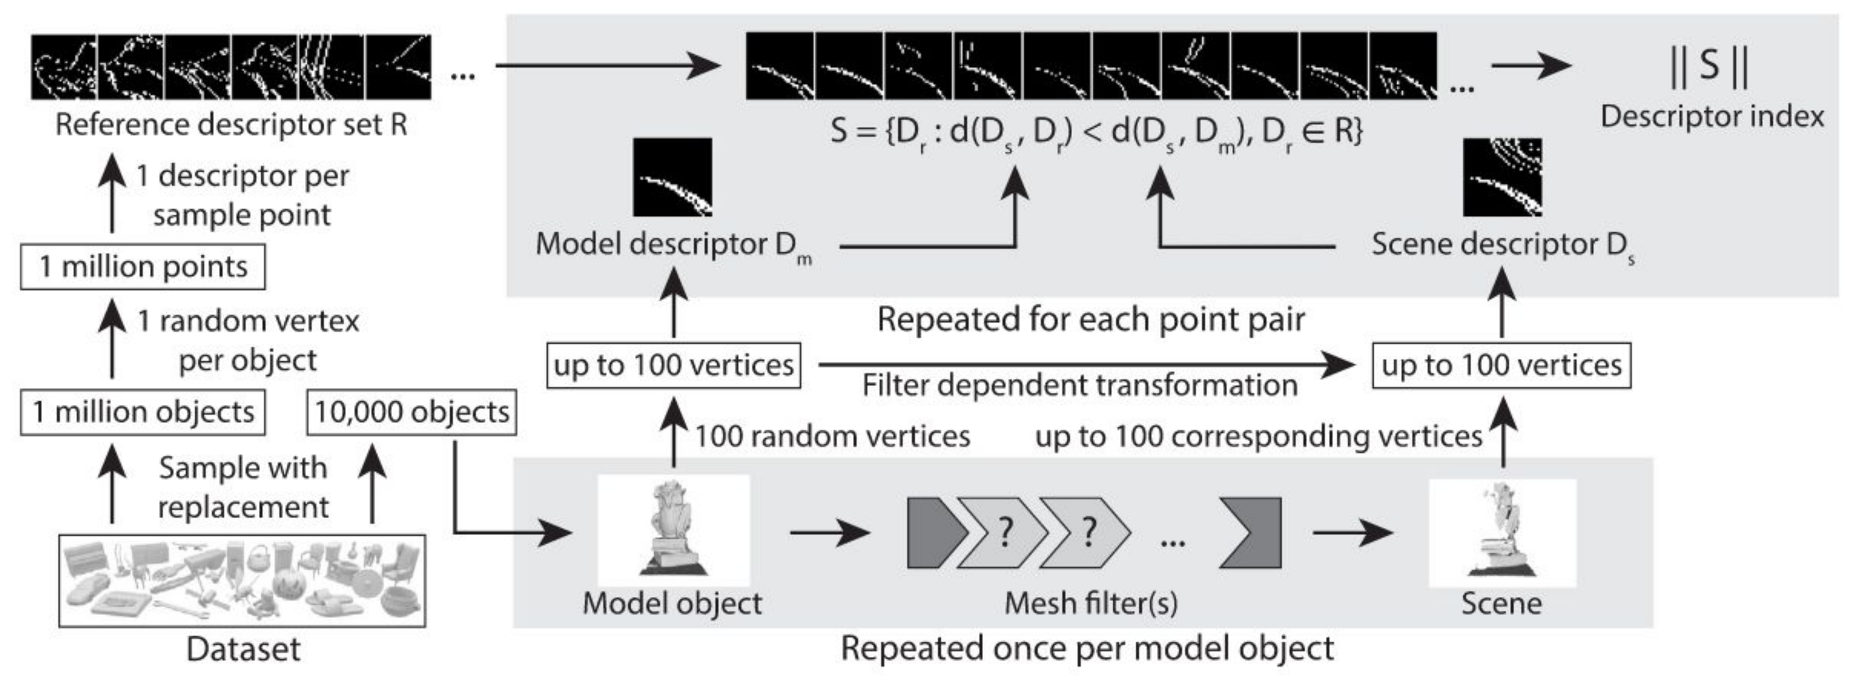

Bart Iver van Blokland, “ShapeBench: A New Approach to Benchmarking Local 3D Shape Descriptors,” Computers & Graphics 124 (November 1, 2024), 4, fig. 2

Results

Occlusion filter results

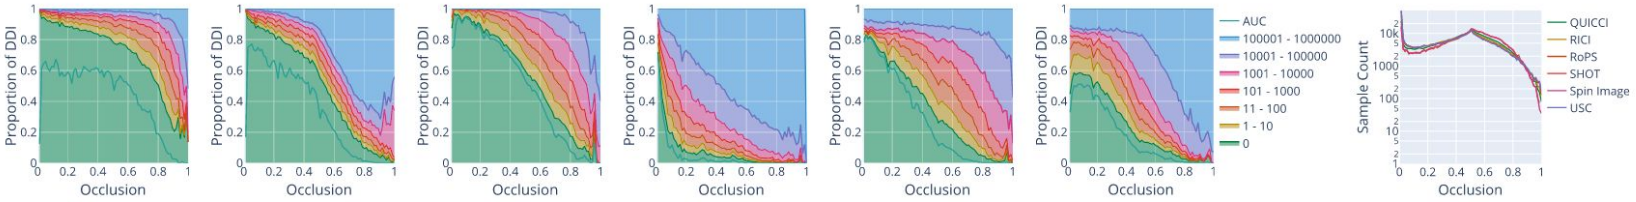

(a) QUICCI (b) RICI (c) Spin Image (d) RoPS (e) SHOT (f) USC (g) Sample distribution

Fig. 8. The effect of varying levels of occlusion on the matching performance of various descriptors. Fig. 8(g) shows the number of sample points per histogram bin.

Bart Iver van Blokland, "ShapeBench: A New Approach to Benchmarking Local 3D Shape Descriptors," Computers & Graphics 124 (November 1, 2024), 8, fig. 8

Single filter results summary

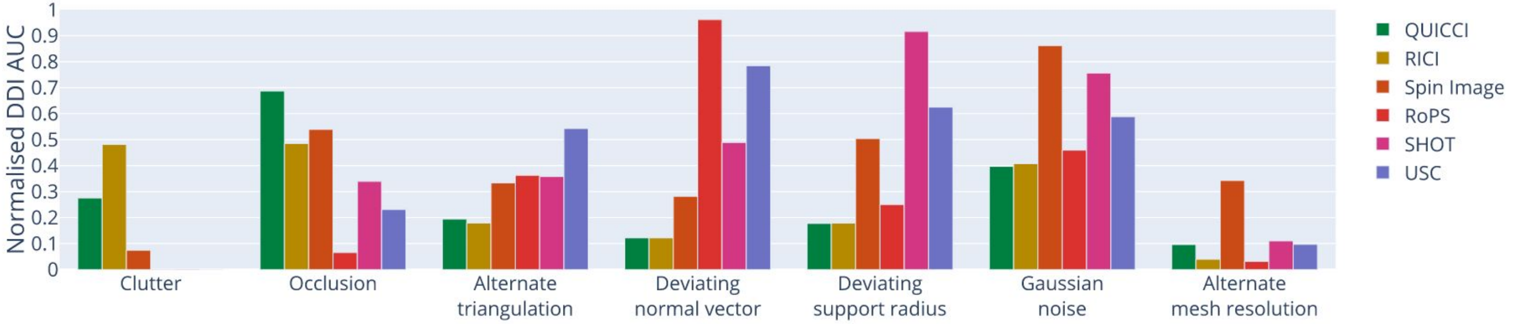

Fig. 14. An overview over the relative performance of the tested methods across each of the filters. Performance is measured as the normalised area under the curve where the proportional DDI is zero.

Bart Iver van Blokland, “ShapeBench: A New Approach to Benchmarking Local 3D Shape Descriptors,” Computers & Graphics 124 (November 1, 2024), 10, fig. 14

Dual filter results

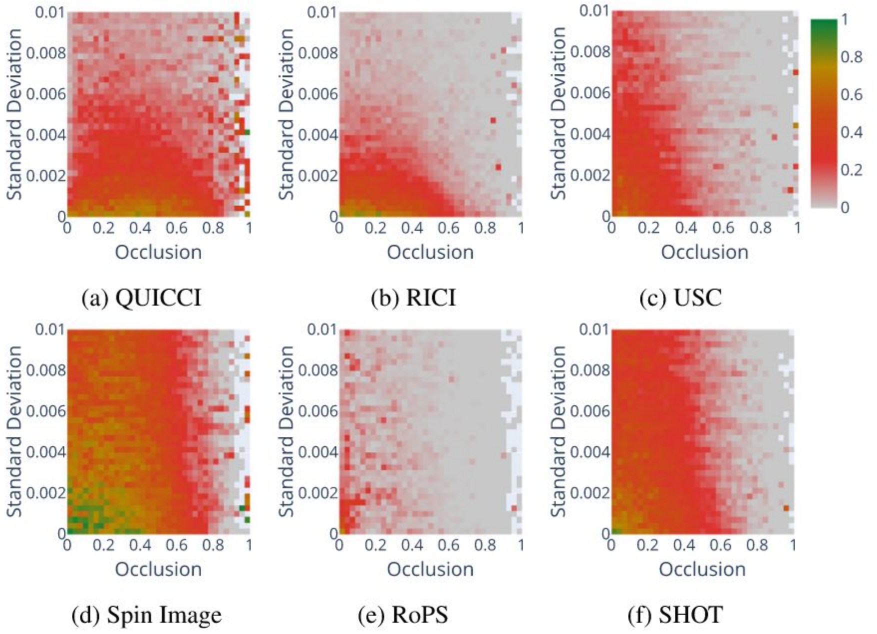

Fig. 17. Results for the occlusion filter followed by the Gaussian noise filter.

Bart Iver van Blokland, “ShapeBench: A New Approach to Benchmarking Local 3D Shape Descriptors,” Computers & Graphics 124 (November 1, 2024), 10, fig. 17

Questions