

Histogram of Distances for Local Surface Description

Odysseas Kechagias-Stamatis, Student *Member, IEEE* and Nabil Aouf

Abstract— 3D object recognition is proven superior compared to its 2D counterpart with numerous implementations, making it a current research topic. Local based proposals specifically, although being quite accurate, they limit their performance on the stability of their local reference frame or axis (LRF/A) on which the descriptors are defined. Additionally, extra processing time is demanded to estimate the LRF for each local patch.

We propose a 3D descriptor which overrides the necessity of a LRF/A reducing dramatically processing time needed. In addition robustness to high levels of noise and non-uniform subsampling is achieved. Our approach, namely Histogram of Distances is based on multiple L2-norm metrics of local patches providing a simple and fast to compute descriptor suitable for time-critical applications. Evaluation on both high and low quality popular point clouds showed its promising performance.

Keywords—3D Object Recognition; 3D Matching; Local Features; Point Cloud

I. INTRODUCTION

Object recognition in 3D is an active research area as it is capable to maintain high performance while being less prone to external conditions such as illumination variation and pose changes [1], [2]. Such properties combined with the commercial low cost 3D information acquiring devices like the “Kinect” have boosted research interest in developing algorithms for object recognition in the 3D domain. Theoretical and practical implementations of such 3D proposals can be found in numerous fields and applications such as robotics, object recognition, surveillance and navigation [3]–[5].

Although several proposals are yet available [4]–[9] all of them demand an accurate and robust Local Reference Frame or Axis (LRF/A) estimation on which the descriptor is calculated. Accuracy and robustness of the LRF/A highly depend on the complexity of the LRF/A algorithm which in turn has a direct impact on the entire processing time. The only exception to that trend is the Heat Kernel Signature [9] with a major disadvantage, i.e. its demand of a great number of RAM memory.

Based on those facts in this paper we propose a solution that overrides the need of a LRF/A reducing dramatically processing time to a much lower level. In parallel the

*Research supported by MBDA UK under grant HP29012014.

O. Kechagias Stamatis and N. Aouf are with the Centre of Electronic Warfare, Cranfield University at the UK Defence Academy, Shrivenham, SN6 8LA, UK (e-mail: {o.kechagiasstamatis, n.aouf}@cranfield.ac.uk)




Fig. 1. Example of matching local descriptors in 3D object recognition scenarios. Top scene from the Bologna dataset which is non-uniformly subsampled to 1/8 its original resolution and Gaussian noise ($\sigma=30\%mr$) is added. Bottom scene is from the SpaceTime dataset. Green lines represent correct matches while red wrong correspondences. Red and blue crosses represent the randomly selected keypoints and their correspondences in respect.

descriptor achieves high performance even under the combination of noise and non-uniform subsampling of the point cloud.

The rest of the paper is organised in the following sections. Section 2 presents a short literature review of the existing 3D pattern recognition algorithms. Section 3 refers to our proposed approach. Section 4 compares and contrasts our approach with the state-of-the-art existing local feature based algorithms. Finally section 5 concludes this paper.

II. RELATED WORK

3D object recognition techniques can broadly be divided into global and local feature based.

Global feature based techniques process and describe the object as one entity and have merely been used in 3D shape retrieval and classification [10]. Although their main advantage is computational efficiency [11], they demand *a priori* segmentation of the target from the scene and are not robust against clutter and occlusion [12]. Examples of Global based techniques are the Shape Distributions [13], VFH [14], CVFH [15], OUR-CVFH [16], ESF [11] and the Projection Density Energy based solution [3].

Local feature based techniques describe local patches around a point of interest providing a valuable solution to partially visible objects in occluded scenes, in object registration, pose estimation and object recognition. Hence, a great number of local feature based pattern recognition

attempts have been made with the trend being an extension of the already mature 2D pattern recognition algorithms to entirely new 3D approaches or solutions based on depth images.

Representatives in the 2D to 3D extension are THRIFT [17], 3D SURF [18], 3D Harris [19] and 3D V-FAST [20]. Their main drawback is that a 3D sensor provides non-volumetric data. Hence a pre-processing step is required to transform the data into voxels which increases the total computational time.

Pure 3D approaches can directly be applied to the point cloud of the target or to its mesh. If mesh information is required, some extra time is needed to calculate the mesh itself, since 3D sensors provide only the relative distance to the target. Among the most well-known mesh based algorithms are RoPS [4], Spin Images [7], HKS [9], TriSI [21] and Tensor [22]. Direct point cloud based solutions are SHOT [5], 3DSC [6], FPFH [8], ISS [23] and the USC [24]. Except for HKS, all the aforementioned solutions demand an accurate Local Reference Frame or Axis (LRF/A) as descriptors are very sensitive to the degree of misalignment between LRF/A's of the corresponding points. It is very challenging though to define a repeatable LRF/A under noise and/or point cloud density variation [25]. This demand adds to those approaches an extra calculation burden which increases the overall processing time. HKS is an exception because it ignores a LRF/A but it demands a great number of RAM memory in the order of 6GB for a point cloud of 30,000 vertices [26].

Depth image pattern recognition is based on 2D projections of a 3D object on a defined reference frame. This hybrid 2D-3D version either exploits 2D methods based on SURF [27] and SIFT [12] or extends to dedicated approaches like BRAND [28] and NARF [29].

The computer vision community has made many positive attempts in 3D pattern recognition achieving high performance while the target is under occlusion and clutter. It is still challenging though to design a descriptor achieving high performance and fast execution time without being prone to high noise levels and point cloud decimation.

Our approach lies on the pure Local feature based 3D approaches and is grounded on three pylons. First, it avoids the LRF/A construction which increases the total processing time and even a minor misalignment affects to a large extent the overall performance. Second, it relies on the simplicity of the descriptor as long as spatial information is maintained. Third, it applies directly the algorithm to the point cloud such as to avoid any additional computational cost.

Inspired by the Shape Distributions [13], the D1 function was modified and applied many times on a Local basis. Specifically, based on a point of interest i.e. keypoint, a spherical volume of radius r was extracted and one point from the border of the local area was randomly selected as a local reference point. From that reference point the L2-norm distances to the vertices belonging to the local area were calculated. The advantage of dealing with distances and not

with angles, as in the majority of the approaches [4], [7], [8], [21], [23], [24], [30], was the fact that the time consuming and sensitive to perturbations LRF/A was not necessary. In addition the relative L2-norm distances of the vertices to the reference point were fully robust to 3D rotations and translations. Our descriptor, named Histogram of Distances (HoD), transformed the distance signature of the local area into a Histogram by compressing information into properly sized and normalised bins. We adopted a normalised Histogram based concept to gain noise invariance and point cloud decimation. Finally, although the reference point was chosen randomly, it had to remain the same throughout the trial.

Even though in the global based approaches the D2 function was proven superior compared to the D1 [11], [13], the calculation of the L2-norm of all point pairs on multiple local spherical volumes increased the total processing time influencing its real-time performance. The proposed HoD descriptor is visualised in Fig. 2 and it showed high descriptive power being robust under heavy noise, point cloud decimation and to their combination.

III. HOD FEATURE DESCRIPTOR

Given a point cloud $P \subset \mathfrak{R}^3$, each point of the cloud can be represented as $P_i = (x_i, y_i, z_i)^T, i \in [0, K]$ where K is the total number of points. For a given set of keypoints, a spherical volume with support radius r centred on each keypoint was extracted. Keypoints could either be randomly selected or by applying the existing 3D keypoint detectors [26]. For each local area, one border point was randomly chosen as a reference point P_r which remained unchanged throughout all trials.

Main difference between our implementation compared to Osada's D1 shape distribution [13] was the extension to a Local basis and the substitution of the reference point by the centroid of the involved points to the border. The latter modification increased the robustness of the HoD descriptor by extending its L2-norm variability. Based on that reference point, the L2-norm was calculated for all vertices in each local area which was properly normalised and discretised to a pre-defined number of bins.

Hence the L2-norm d_i between the reference point P_r and each point $P_i, i \in [0, L], L \leq K$ of the local area was given by:

$$d_i = \|P_r - P_i\|_2 \quad (1)$$

Noise invariance was achieved via the normalisation of the d_i distances and by obtaining the round-floor values d_i' :

$$d_i' = \left\lfloor \frac{d_i}{\max(d_i)} B \right\rfloor \quad (2)$$

Thereafter, the normalised distances d_i' were encoded into a histogram namely the Histogram of Distances (HoD). To increase the descriptiveness of HoD, we concatenated the




Fig. 2. Histogram of Distances (HoD) concept. (a) A spherical area of radius r (red) centred on a keypoint is extracted. (b) One random border point from the local area is selected as reference point (yellow) and the reference point to vertices L2-norms distances are calculated (in red as example). (c) The coarse and fine normalised distances were encoded into a Histogram of Distances

coarse and fine normalised distances d_i' by selecting properly sized bins B . Point cloud resolution invariance was accomplished by normalising the HoD descriptor to sum up to 1 [5].

Existing approaches define the support radius of the local area as a multiple of the average model mesh resolution (\overline{mr}). We extended the mesh resolution invariance by substituting the support radius metric with a multiple of each scene's mesh resolution (mr). Although \overline{mr} was equal to the average resolution of all scenes, each individual mr had a minor fluctuation from scene to scene. This methodology assisted much during the varying mesh resolution trials as each model's local area support radius was adapted to each scene. A minor drawback was the slight extension of the processing burden as mr had to be real-time calculated.

IV. EXPERIMENTAL RESULTS

Evaluation of the HoD feature descriptor was performed on the Bologna and SpaceTime dataset [5] and was compared with state-of-the-art feature descriptors.

A. Experimental Setup on the Bologna dataset

1) The Bologna dataset was exploited comprising of 6 models and 45 scenes. Models were taken from the Stanford 3D Scanning Repository [31] and were randomly rotated and translated to create clutter and pose variations. In contrast to [32] we exploited the entire dataset and not a subset.

2) The criterion to assess the descriptor was the popular 1-precision vs. recall curve (PR) as in most recent papers [17], [21], [28], [30], [32]–[34]. The PR curve was based on a set of model features f_i^M , a ground truth transformation and the corresponding scene features f_i^S . A scene feature was matched with all model features based on their Euclidean distance and the Nearest Neighbour Distance Ratio criterion. If the ratio of the nearest model feature f_i^M with the second

nearest f_i^M were less than a threshold τ , then the scene feature f_i^S and the model feature f_i^M were considered as a match. Furthermore, if the Euclidean distance of the physical location of the matched keypoints was less than half the HoD's descriptor support radius, then the match was considered a True Positive (TP) and otherwise as a False Positive (FP). For fairness reasons we used the HoD's support radius for all descriptors regardless of their individual values. Correspondences were established in the same manner. Recall and 1-precision were defined as [34]:

$$recall = \frac{\#TP}{\#correspondences} \quad (3)$$

$$1 - precision = \frac{\#FP}{\#matches} \quad (4)$$

By altering the NNDR threshold values (τ) in the range [0,1] we obtained the PR curve which ideally would be at the upper left corner.

3) We compared the HoD descriptor with the state-of-the-art descriptors SHOT [5], RoPS [4], and FPFH [8]. Each descriptor's support radius was independently tuned on training scenes similar to the Bologna dataset. These scenes were non-uniformly subsampled to $\frac{1}{2}$ their mesh resolution and Gaussian noise was added with a standard deviation of 10% the average mesh resolution \overline{mr} [5], [35]. In contrast to [30], [32], [36] we chose the non-uniform sampling as in reality laser beam distortions can influence both the total number of points the cloud consists of as well as their spatial location.

Time critical applications were our main implementation interest, therefore we randomly selected 100 keypoints from each model and extracted their corresponding ones in the scene based on their *a priori* known ground truth transformation. Random keypoint selection was preferred against exploiting a keypoint detector as errors of the

detector did not affect the descriptor. Approaches listed in table 1 were used to extract the local features of all keypoints both for the model and their corresponding ones in the scenes. Performance evaluation was based on their PR curve.

All trials and benchmarks were performed in MATLAB 2015a and C++. Implementations in C++ were obtained from the Point Cloud Library (PCL) Version 1.7.2 [37] and RoPS from MATLAB File Exchange [38]. Beyond the support radius which was equal to the HoD's, the rest of the parameters were fixed either to the ones originally proposed by their authors or to their PCL implementation [39]. Although FPFH had the smallest support radius compared to the rest of the descriptors, its performance peaked at a value of $20\overline{mr}$. We confirm the evaluation in [39] which states that FPFH performance peaks for a certain support radius value and beyond that its performance drops as well as [40] regarding the direct relationship of SHOT with its support radius.

The distance match metric for the HoD feature descriptor was the Euclidean while the rest feature descriptors used the original metric proposed by their authors. All approaches benefit from the kd-tree indexing in order to speed-up the matching stage.

B. Robustness to noise

We evaluated and compared the robustness of our HoD descriptor against the descriptors of table 1 under various heavy noise levels. Trials included the addition of Gaussian noise with $\sigma = \{200\%, 300\%\}\overline{mr}$ to the 45 scenes of the Bologna dataset and was independently added to the x, y and z axes of each scene point [33]. For each noise level the PR curve generated is presented in Fig. 3.

HoD and RoPS achieved best performance compared to SHOT and FPFH. For $\sigma = 200\%\overline{mr}$ case, RoPS achieved a slightly higher recall compared to HoD. In the case of severe noise level though with $\sigma = 300%\overline{mr}$, HoD outperformed RoPS. In both noise scenarios SHOT achieved moderate performance while FPFH was very sensitive to such a high noise level confirming [32].

In any case the computational cost of HoD was much lower compared to RoPS with a detailed processing time analysis following in Section IV-E.

HoD achieved robustness to noise due to at least three factors. First, by ignoring the sensitive to noise LRF estimation on which the rest of the descriptors rely. Second, by using a border point as a local reference point. Instead of the fixed centroid, as proposed by Osada [13], our reference point was equally affected by noise like all the vertices in the local area, so distance perturbation due to noise had a minor impact. Third, the descriptor bin B was partially large sized so that distance fluctuations due to noise still fall into the original noise-free bins.

Table I: Parameter values for the descriptors

	Support radius	Descriptor Length	Implementation
RoPS	$40\overline{mr}$	135	MATLAB
SHOT	$40\overline{mr}$	352	C++ (PCL)
FPFH	$20\overline{mr}$	33	C++ (PCL)
HoD	$40\overline{mr}$	240	MATLAB

C. Robustness to Varying Mesh Resolutions

We evaluated and compared the robustness of our HoD descriptor against the descriptors of table 1 under various subsampled resolutions. In contrast to [30], [32], [36] we non-uniformly subsampled the noise-free scenes to $\{1/4, 1/8\}$ of their original mesh resolution. For each noise level the PR curve generated is presented in Fig. 4.

HoD outperformed all competitors in both subsampling cases followed closely by RoPS. In both sub-sampling scenarios HoD gained both highest recall and precision. Although in the 1/4 subsampling case SHOT and FPFH had similar performance, in the 1/8 non-uniform subsampling case SHOT performed slightly better.

D. Robustness to combined Gaussian Noise and Varying Mesh Resolutions

Our HoD descriptor was evaluated and compared against the descriptors of table 1 under various combinations of non-uniform subsampling and Gaussian noise. Specifically, trials included 1/2 subsampling with $\sigma = 10\%\overline{mr}$ and 1/8 with $\sigma = 30\%\overline{mr}$. For each subsampling - noise level combination the PR curve generated is presented in Fig. 5.

For the first case, with non-uniform subsampling to 1/2 the original resolution and with Gaussian noise with $\sigma = 10\%\overline{mr}$, HoD outperformed all approaches showing a robust performance under combined noise and point cloud subsampling.

Regarding the 1/8 subsampling with $\sigma = 0.3\overline{mr}$ noise case, HoD and RoPS achieved equally the highest performance. SHOT and FPFH achieved similar recall with SHOT though having greater precision.

It is worth mentioning that in any noise level - subsampling case HoD was much more processing efficient compared to RoPS.

E. Processing time

For a fair comparison we compared the HoD descriptor only with RoPS as both were implemented in MATLAB. Comparison with SHOT and FPFH which were implemented in C++ is not valid as by definition C++ is faster. But even in that case, HoD in MATLAB was faster than FPFH in C++ due to the size of the support radius of the local area.

Even though HoD included fully real-time point

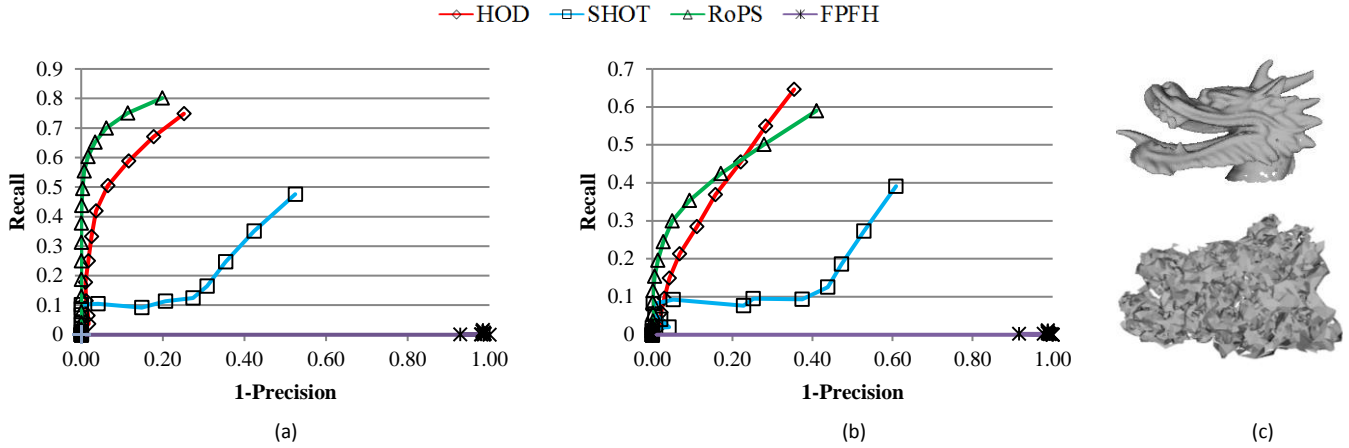


Fig. 3. PR curves under various Gaussian noise levels (a) $\sigma=200\% mr$ (b) $\sigma=300\% mr$ (c) Original and with $\sigma=200\% mr$ Gaussian noise (in mesh representation for better visibility)

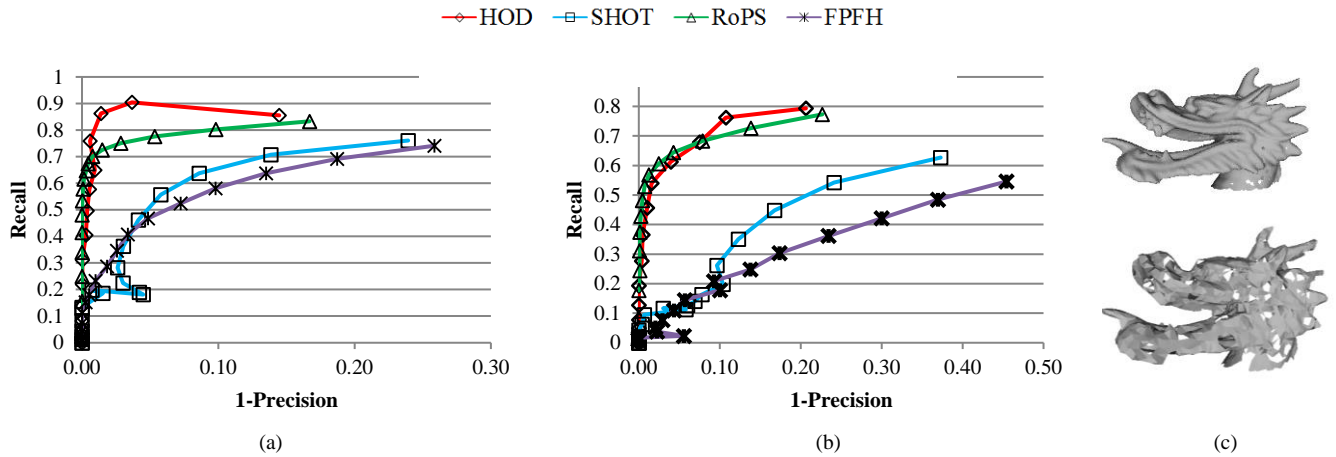


Fig. 4. PR curves under varying mesh resolution (a) $1/4$ (b) $1/8$ (c) Original model and $1/8$ Non-Uniform Subsampling (in mesh representation for better visibility)

resolution estimation, neglecting the LRF calculation we reduced dramatically the processing time. Specifically HoD was 7.8 times faster than RoPS and 7 times faster than FPFH. HoD demanded 16ms/correspondence making our solution applicable to real-time applications even when it was implemented in MATLAB.

F. Experimental Setup on the SpaceTime stereo dataset

We further evaluated the HoD descriptor on the SpaceTime dataset [30] which consisted of 6 models and 11 scenes. Compared with the previously tested Bologna dataset, the SpaceTime was harder as it included models and scenes with fewer details. Trials considered the noise-free case using the same local area support radius while texture information was not taken into account. Due to the lower quality data, all descriptors performed poorer than previously.

Higher recall and precision was achieved by HoD, followed by SHOT and RoPS as shown in Fig. 6. In terms of processing efficiency under the MATLAB platform, HoD was 7.2 times faster than RoPS.




Fig. 6. PR curves on the SpaceTime dataset

V. CONCLUSION

We presented a simple and fast-to-calculate 3D descriptor, named the Histogram of Distances (HoD). By overriding the necessity of a Local Reference Frame or Axis estimation, we gained an overall processing time speedup while we maintained robustness to noise and/or point cloud resolution variation under clutter and occlusion. Compared to state-of-the-art approaches on two popular high and low resolution datasets, HoD gained a higher performance in total. HoD's low processing time makes it a solution to time-critical applications in the field of computer vision based scenarios.




Fig. 5. PR curves under combined varying mesh resolution and Gaussian noise (a) $\frac{1}{2}$ & $\sigma=10\% mr$ (magnified) (b) $\frac{1}{8}$ & $\sigma=30\% mr$ (c) $\frac{1}{2}$ Non-Uniform subsampled with $10\% mr$ noise (top) $\frac{1}{8}$ Non-Uniform subsampled with $30\% mr$ noise (bottom) in mesh representation for better visibility

REFERENCES

- [1] A. S. Mian, M. Bennamoun, and R. Owens, "Three-dimensional model-based object recognition and segmentation in cluttered scenes," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 28, no. 10, pp. 1584–601, Oct. 2006.
- [2] Y. Lei, H. Lai, and X. Jiang, "3D face recognition by SURF operator based on depth image," in *Proceedings - 2010 3rd IEEE International Conference on Computer Science and Information Technology, ICCSIT 2010*, 2010, vol. 9, pp. 240–244.
- [3] O. Kechagias-Stamatis and N. Aouf, "Fast 3D object matching with Projection Density Energy," in *2015 23rd Mediterranean Conference on Control and Automation (MED)*, 2015, pp. 752–758.
- [4] Y. Guo, F. Sohel, M. Bennamoun, M. Lu, and J. Wan, "Rotational Projection Statistics for 3D Local Surface Description and Object Recognition," *Int. J. Comput. Vis.*, vol. 105, no. 1, pp. 63–86, Apr. 2013.
- [5] F. Tombari, S. Salti, and L. Di Stefano, "Unique signatures of histograms for local surface description," *Comput. Vision—ECCV 2010*, pp. 356–369, 2010.
- [6] A. Frome, D. Huber, R. Kolluri, T. Bülow, and J. Malik, "Recognizing Objects in Range Data Using Regional Point Descriptors," in *European Conference on Computer Vision*, 2004, pp. 224–237.
- [7] A. E. Johnson and M. Hebert, "Using spin images for efficient object recognition in cluttered 3D scenes," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 21, no. 5, pp. 433–449, 1999.
- [8] R. B. Rusu, N. Blodow, and M. Beetz, "Fast Point Feature Histograms (FPFH) for 3D registration," *2009 IEEE Int. Conf. Robot. Autom.*, pp. 3212–3217, 2009.
- [9] J. Sun, M. Ovsjanikov, and L. Guibas, "A concise and provably informative multi-scale signature based on heat diffusion," *Eurographics Symp. Geom. Process.*, vol. 28, no. 5, pp. 1383–1392, 2009.
- [10] B. Bustos, D. Keim, D. Saupe, T. Schreck, and D. V. Vranić, "Feature-based similarity search in 3D object databases," *ACM Comput. Surv.*, vol. 37, no. 4, pp. 345–387, 2005.
- [11] W. Wohlkinger and M. Vincze, "Ensemble of shape functions for 3D object classification," in *2011 IEEE International Conference on Robotics and Biomimetics*, 2011, pp. 2987–2992.
- [12] N. Bayramoglu and A. A. Alatan, "Shape Index SIFT: Range Image Recognition Using Local Features," in *2010 20th International Conference on Pattern Recognition*, 2010, pp. 352–355.
- [13] R. Osada, T. Funkhouser, B. Chazelle, and D. Dobkin, "Matching 3D models with Shape distribution," in *Shape Modeling and Applications, SMI 2001 International Conference*, 2001, pp. 154–166.
- [14] R. B. Rusu, G. Bradski, R. Thibaux, and J. Hsu, "Fast 3D recognition and pose using the Viewpoint Feature Histogram," *2010 IEEE/RSJ Int. Conf. Intell. Robot. Syst.*, pp. 2155–2162, Oct. 2010.
- [15] A. Aldoma, M. Vincze, N. Blodow, D. Gossow, S. Gedikli, R. B. Rusu, and G. Bradski, "CAD-model recognition and 6DOF pose estimation using 3D cues," in *2011 IEEE International Conference on Computer Vision Workshops (ICCV Workshops)*, 2011, pp. 585–592.
- [16] A. Aldoma, F. Tombari, R. B. Rusu, and M. Vincze, "OUR-CVFH - Oriented, unique and repeatable clustered viewpoint feature histogram for object recognition and 6DOF pose estimation," *Lect. Notes Comput. Sci. (including Subser. Lect. Notes Artif. Intell. Lect. Notes Bioinformatics)*, vol. 7476 LNCS, pp. 113–122, 2012.
- [17] A. Flint, A. Dick, and A. Van Den Hengel, "Thrifty: Local 3D Structure Recognition," *9th Bienn. Conf. Aust. Pattern Recognit. Soc. Digit. Image Comput. Tech. Appl. (DICTA 2007)*, pp. 182–188, Dec. 2007.
- [18] J. Knopp, M. Prasad, G. Willems, R. Timofte, and L. Van Gool, "Hough transform and 3D SURF for robust three dimensional classification," *Comput. Vis. – ECCV 2010 Lect. Notes Comput. Vis.*, vol. 6316, pp. 589–602, 2010.
- [19] I. Sipiran and B. Bustos, "Harris 3D: a robust extension of the Harris operator for interest point detection on 3D meshes," *Vis. Comput.*, vol. 27, no. 11, pp. 963–976, Jul. 2011.
- [20] S. Koelstra and I. Patras, "The fast-3D spatio-temporal interest region detector," in *2009 10th Workshop on Image Analysis for Multimedia Interactive Services*, 2009, pp. 242–245.
- [21] Y. Guo, F. Sohel, M. Bennamoun, M. Lu, and J. Wan, "TriSI: A Distinctive Local Surface Descriptor for 3D Modeling and Object Recognition," in *8th International Conference on Computer Graphics Theory and Applications*, 2013.
- [22] A. S. Mian, M. Bennamoun, and R. A. Owens, "A Novel Representation and Feature Matching Algorithm for Automatic Pairwise Registration of Range Images," *Int. J. Comput. Vis.*, vol. 66, no. 1, pp. 19–40, Jan. 2006.
- [23] Y. Zhong, "Intrinsic shape signatures: A shape descriptor for 3D object recognition," *2009 IEEE 12th Int. Conf. Comput. Vis. Work. ICCV Work.*, pp. 689–696, Sep. 2009.
- [24] F. Tombari, S. Salti, and L. Di Stefano, "Unique Shape Context for 3D Data Description," in *Proceedings of the ACM Workshop on 3D Object Retrieval*, 2010, pp. 57–62.
- [25] A. Petrelli and L. Di Stefano, "On the repeatability of the local reference frame for partial shape matching," *2011 Int. Conf. Comput. Vis.*, pp. 2244–2251, Nov. 2011.
- [26] F. Tombari, S. Salti, and L. Di Stefano, "Performance evaluation of 3D keypoint detectors," *Int. J. Comput. Vis.*, vol. 102, no. 1–3, pp. 198–220, 2013.

DOI: 10.1109/ICRA.2016.7487402

- [27] T.-W. R. Lo and J. P. Siebert, "Local feature extraction and matching on range images: 2.5D SIFT," *Comput. Vis. Image Underst.*, vol. 113, no. 12, pp. 1235–1250, Dec. 2009.
- [28] E. R. Nascimento, G. L. Oliveira, M. F. M. Campos, A. W. Vieira, and W. R. Schwartz, "BRAND: A robust appearance and depth descriptor for RGB-D images," in *2012 IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2012, pp. 1720–1726.
- [29] S. Bastian, R. R. Bogdan, K. Kurt, and W. Burgard, "NARF: 3D Range Image Features for Object Recognition," *Work. Defn. Solving Realis. Percept. Probl. Pers. Robot. IEEE/RSJ Int. Conf. Intell. Robot. Syst.*, vol. 215, no. 1, p. 11, 2010.
- [30] S. Salti, F. Tombari, and L. Di Stefano, "SHOT: Unique signatures of histograms for surface and texture description," *Comput. Vis. Image Underst.*, vol. 125, pp. 251–264, 2014.
- [31] B. Curless and M. Levoy, "A Volumetric Method for Building Complex Models from Range Images."
- [32] Y. Guo, M. Bennamoun, F. Sohel, M. Lu, J. Wan, and J. Zhang, "Performance Evaluation of 3D Local Feature Descriptors," in *Computer Vision -- ACCV 2014*, D. Cremers, I. Reid, H. Saito, and M.-H. Yang, Eds. 2015, pp. 178–194.
- [33] Y. Guo, F. a. Sohel, M. Bennamoun, J. Wan, and M. Lu, "RoPS: A local feature descriptor for 3D rigid objects based on rotational projection statistics," *2013 1st Int. Conf. Commun. Signal Process. their Appl.*, pp. 1–6, Feb. 2013.
- [34] K. Mikolajczyk and C. Schmid, "Performance evaluation of local descriptors," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 27, no. 10, pp. 1615–30, Oct. 2005.
- [35] Y. Guo, M. Bennamoun, F. Sohel, M. Lu, and J. Wan, "3D Object Recognition in Cluttered Scenes with Local Surface Features: A Survey," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 36, no. 11, pp. 2270–2287, Nov. 2014.
- [36] Y. Guo, M. Bennamoun, F. A. Sohel, J. Wan, and M. Lu, "3D free form object recognition using rotational projection statistics," in *2013 IEEE Workshop on Applications of Computer Vision (WACV)*, 2013, pp. 1–8.
- [37] R. B. Rusu and S. Cousins, "3D is here: Point Cloud Library (PCL)," in *2011 IEEE International Conference on Robotics and Automation*, 2011, pp. 1–4.
- [38] "MATLAB Central FileExchange." [Online]. Available: <http://www.mathworks.com/matlabcentral/fileexchange/>. [Accessed: 20-Apr-2014].
- [39] Y. Guo, M. Bennamoun, F. Sohel, M. Lu, J. Wan, and N. Ming, "A Comprehensive Performance Evaluation of 3D Local Feature Descriptors."
- [40] M. Restrepo and J. Mundy, "An Evaluation of Local Shape Descriptors in Probabilistic Volumetric Scenes," in *Proceedings of the British Machine Vision Conference 2012*, 2012, pp. 46.1–46.11.