

Signature of Geometric Centroids for 3D Local Shape Description and Partial Shape Matching

Keke Tang, Peng Song^(✉) and Xiaoping Chen

University of Science and Technology of China

kktang@mail.ustc.edu.cn, songpeng@ustc.edu.cn, xpchen@ustc.edu.cn

Abstract. Depth scans acquired from different views may contain nuisances such as noise, occlusion, and varying point density. We propose a novel *Signature of Geometric Centroids* descriptor, supporting direct shape matching on the scans, without requiring any preprocessing such as scan denoising or converting into a mesh. First, we construct the descriptor by voxelizing the local shape within a uniquely defined local reference frame and concatenating geometric centroid and point density features extracted from each voxel. Second, we compare two descriptors by employing only corresponding voxels that are both non-empty, thus supporting matching incomplete local shape such as those close to scan boundary. Third, we propose a descriptor saliency measure and compute it from a descriptor-graph to improve shape matching performance. We demonstrate the descriptor’s robustness and effectiveness for shape matching by comparing it with three state-of-the-art descriptors, and applying it to object/scene reconstruction and 3D object recognition.

1 Introduction

The recent development in depth sensing devices offers a convenient and flexible way to acquire depth scans of an object or a scene that represent their partial shapes. In practice, we need to register these scans into a common coordinate system to better understand the object’s or scene’s geometry [1] or compare known object models with these scans for 3D object recognition [2]. All these applications require solving the partial shape matching problem [3, 4].

Depth scans (i.e., 3D point clouds) lack topology information of the shape and usually contain noise, holes, and/or varying point density. To facilitate partial shape matching, one common way is to convert the point cloud into a mesh to remove the noise and fill the holes, and then perform shape matching on the mesh instead [5–8]. Although this conversion simplifies the matching process, it brings several drawbacks. First, original partial shape could be modified and/or downsampled by the conversion, e.g., when smoothing the depth scan for denoising. Second, the mesh topology generated by the conversion could be different from the real one such as incorrectly filled holes, misleading the shape matching.

Therefore, other researchers seek to perform shape matching directly on the point cloud data. This is generally achieved by representing and matching the scans using local shape descriptors. Although existing descriptors [9–12] work

well on clean depth scans, they have difficulties dealing with original scans acquired under various conditions such as occlusion, clutter, and varying lighting. This is because these descriptors are sensitive to noise and/or varying point density due to their encoded shape features such as point density [9, 10] and surface normals [11], or are sensitive to scan boundary and holes due to their descriptor comparison scheme that is based on the vector distance [11, 12].

To address above limitations, we propose a *Signature of Geometric Centroids* (SGC) descriptor for partial shape matching with three novel components:

- *A Robust Descriptor.* We construct the SGC descriptor by voxelizing the local shape within a uniquely defined local reference frame (LRF) and concatenating the geometric centroid and point density features extracted from each non-empty voxel. Thanks to the extracted shape features, our descriptor is robust against noise and varying point density.
- *A Descriptor Comparison Scheme.* Rather than simply computing the Euclidean distance between two descriptors, we compute a similarity score between two descriptors based on comparing the extracted features from corresponding voxels that are both non-empty. By this, the comparison scheme supports shape matching between local shape that are incomplete.
- *Descriptor Saliency for Shape Matching.* Different from keypoint detection [13] that identifies distinct points locally on a single scan/model, we propose descriptor saliency to measure distinctiveness of SGC descriptors across all input scans and compute it from a descriptor-graph. Guided by the descriptor saliency, we improve shape matching performance by intentionally selecting distinct descriptors to find corresponding feature points.

We evaluate the robustness of SGC against various nuisances including scan noise, varying point density, distance to scan boundary, occlusion, and the effectiveness of using SGC and descriptor saliency for partial shape matching. Experimental results show that SGC outperforms three start-of-the-art descriptors (i.e., spin image [9], 3D shape context [10], and signature of histograms of orientations (SHOT) [11]) on publicly available datasets. We further apply SGC to two typical applications of partial shape matching, i.e., object/scene reconstruction and 3D object recognition, to demonstrate its usefulness in practice.

2 Related Work

Shape Matching. Shape matching aims at finding correspondences between complete or partial models by comparing their geometries. Many shape matching approaches apply global shape descriptors to characterize the whole shape, for example, using Reeb graphs [14] or skeleton graphs [15] for articulated objects and shape distributions [16] for rigid objects. However, depth scans acquired from each single view usually have significant missing data. Matching these partial shapes is a difficult task because, before computing the correspondences of the shapes, we first need to find the common portions among them [1]. This requires a careful design of local shape descriptors [17] that are less sensitive to occlusion.

Local Shape Descriptors. Local shape descriptors can be classified as low- and high-dimensional, according to the richness of encoded local shape information. Low-dimensional descriptors such as surface curvature [18] and surface hashes [19], are easy to compute, store, and compare, yet have limited descriptive ability. Compared with them, high-dimensional descriptors provide a fairly detailed description of the local shape around a surface point. We classify high-dimensional descriptors into three classes according to their attached LRF [20].

Descriptors without an LRF. Early local shape descriptors are generated by directly accumulating some geometric attributes into a histogram, without building an LRF. Hetzel et al. [21] represented local shape patches by encoding three local shape features (i.e., pixel depth, surface normals, and curvatures) into a multi-dimensional histogram. Yamany et al. [22] described local shape around a feature point by generating a signature image that captures surface curvatures seen from that point. Kokkinos et al. [23] generated an intrinsic shape context descriptor by shooting geodesic outwards from a keypoint to chart the local surface and creating a 2D histogram of features defined on the chart.

Due to the missing of an LRF, the correspondence built by matching the descriptors is limited to the point spatial position only. Thus, to match two scans by estimating a rigid transform, at least three pairs of corresponding points need to be found, making the space of searching corresponding points large.

Descriptors with a non-unique LRF. Researchers later attached an LRF for local shape descriptors to enrich the correspondence with spatial orientation. By this, two scans can be matched by finding a single pair of corresponding points using the descriptors and estimating the transform based on aligning associated LRFs. However, since the attached LRF is not unique, a further disambiguation process is required for the generated transform.

Johnson et al. [24] proposed a spin image descriptor by spinning a 2D image about the normal of a feature point and summing up the number of points that fall into the bins of that image. Frome et al. [10] proposed a 3D shape context (3DSC) descriptor by generating a 3D histogram of accumulated points within a partitioned spherical volume centered at a feature point and aligned with the feature normal. Mian et al. [5] proposed a 3D tensor descriptor by constructing an LRF from a pair of oriented points and encoding the intersected surface area into a multidimensional table. Zhong [25] proposed intrinsic shape signatures by improving [10] based on a different partitioning of the 3D spherical volume and a new definition of LRF with ambiguity.

Descriptors with a unique LRF. Recently, researchers constructed a unique LRF from the local shape around a feature point and further describe the local shape relative to the LRF. Thanks to the unique LRF, the transform to match two scans can be uniquely defined based on aligning corresponding LRFs.

Tombari et al. [11] proposed a SHOT descriptor by concatenating local histograms of surface normals defined on each bin of a partitioned spherical volume aligned with a unique LRF. Guo et al. [7] constructed a RoPS descriptor by rotationally projecting the neighboring points of a feature point onto 2D planes and calculating a set of statistics within a unique LRF. Guo et al. [12] later gen-

erated three signatures representing the point distribution in three cylindrical coordinate systems and concatenated and compressed these signatures into a Tri-Spin-Image descriptor. Song and Chen [8] developed a local voxelizer descriptor by voxelizing local shape within a unique LRF and concatenating an intersected surface area feature in each voxel, and applied it to surface registration [26].

SGC is also constructed within a unique LRF. Compared with above descriptors, the geometric centroid feature that we extract for constructing the descriptor is more robust against noise and varying point density. Moreover, our descriptor comparison scheme supports matching local shape that is close to the scan boundary. By this, SGC is more robust for shape matching on point cloud data than state-of-the-art descriptors [9–11], see Section 5 for the comparisons.

3 Signature of Geometric Centroids Descriptor

This section presents the method to construct an SGC descriptor for the local shape (i.e., support) around a feature point p , a scheme to compare a pair of SGC descriptors, and the parameters tuned for generating SGC descriptors.

3.1 LRF Construction

Given a feature point p on a scan and a radius r , a local support is defined by intersecting the scan with a sphere centered at p with radius r . Taking this support as input, we construct a unique LRF based on principal component analysis (PCA) on the support by using the approach in [11], see Figure 1(a). When the normal of p is available, we further improve the disambiguation of LRF axes by enforcing the principal axis associated with the smallest eigenvalue (i.e., the blue axis in Figure 1(a)) to be consistent with the normal [8].

3.2 SGC Construction

Given the unique LRF, a general way to construct a descriptor is to partition a support into bins, extract shape features from each bin, and concatenate the values representing the shape features into a descriptor vector (or a histogram). **Partition the Support.** Given a support S_p around a feature point p , there are three typical approaches to partition S_p into small local patches. The first one is to partition the bounding spherical volume of S_p into girds evenly [11] or logarithmically [10] along azimuth, elevation and radial dimensions. The second one is to partition the angular space of the spherical volume into relatively homogeneously distributed bins [25]. However, the bins generated by these two approaches have varying sizes, which need to be compensated when constructing a descriptor. In addition, the irregular shape of these bins complicates the segmentation of local shape within each bin for extracting local shape features.

The third approach is to construct a bounding cubical volume of S_p that is aligned with the LRF and partition the cubical volume into regular bins (i.e., voxels) [8]. These regular bins simplify the extraction of local shape features and




Fig. 1. Constructing an SGC descriptor. (a) Construct a unique LRF from a spherical support centered at a feature point (in pink); (b) segment a cubical support centered at the feature point and aligned with the LRF; (c) voxelize the support and extract centroid features from non-empty voxels; the centroid color indicates point density in the voxel, where small and large densities are colored in blue and red respectively.

thus the descriptor construction. Therefore, we employ the third approach to partition S_p for constructing the SGC descriptor, see Figure 1(b&c). Note that the edges of the cubical volume have a length of $2R$, where $R \geq r$.

Extract Bin Features. Due to the missing of topology information, point clouds have limited types of shape features that can be extracted, e.g., surface normal feature in SHOT [11] and point density feature in 3DSC [10]. This paper proposes extracting a geometric centroid feature from each non-empty voxel for constructing SGC due to following reasons. First, centroid is an integral feature [27], thus can be more robust against noise and varying point density. Second, centroid can be computed simply by averaging the positions of all points staying within a voxel. Note that we do not realize any existing work that employs centroid features for constructing a usable descriptor.

Construct the Descriptor. We divide the cubical volume evenly into $K \times K \times K$ bins (i.e., voxels) with the same size, see Figure 1(c). For each voxel V_i , we identify all N_i points staying within the voxel and then calculate the centroid (X_i, Y_i, Z_i) for the points. Note that, the position of the centroid is relative to the minimum corner of V_i in the LRF. We save the extracted feature as (X_i, Y_i, Z_i, N_i) for non-empty voxels, and $(0,0,0,0)$ for empty ones. An SGC descriptor is generated by concatenating all these values assigned for each voxel. The dimension of an SGC descriptor saved in this way is $4 \times K \times K \times K$.

Thanks to the unique LRF, the three positional values of V_i 's centroid (X_i, Y_i, Z_i) can be compressed into a single value using $C_i = (Z_i \times L + Y_i) \times L + X_i$, where $L = 2R$ denotes the edge length of V_i . By this, we compress the dimension of the descriptor to $2 \times K \times K \times K$, saving 50% storage space.

3.3 Comparing SGC Descriptors

Ideally, SGC descriptors generated for two corresponding points in different scans should be exactly the same. However, due to variance of sampling, noise and occlusion, the two descriptors usually have a certain amount of difference. Unlike existing approaches that compare descriptors by computing their Euclidean distance [11, 7, 8], we develop a new scheme for comparing two SGC descriptors.

When constructing an SGC descriptor, most of the voxels are likely to be empty (see again Figure 1(c)). We classify each pair of corresponding voxels into three cases: 1) empty voxel vs empty voxel; 2) non-empty voxel vs empty voxel; and 3) non-empty voxel vs non-empty voxel. In all three cases, only case 3 should contribute to computing a similarity score between two descriptors. Thus, to compare two SGC descriptors quantitatively, we propose to accumulate a similarity score for every pair of corresponding voxels that are both non-empty.

In detail, we denote two SGC descriptors as D_m and D_n . The similarity between the i -th voxel of D_m , V_m^i , and the i -th voxel of D_n , V_n^i , is defined as:

$$s(V_m^i, V_n^i) = \begin{cases} \ln \frac{N_m^i N_n^i}{\|C_m^i - C_n^i\|^2 + \epsilon}, & \text{for } N_m^i > 0 \text{ and } N_n^i > 0 \\ 0 & \text{for } N_m^i = 0 \text{ or } N_n^i = 0 \end{cases} \quad (1)$$

where N_m^i and N_n^i represent the number of points in V_m^i and V_n^i respectively, while C_m^i and C_n^i represent the centroid of V_m^i and V_n^i respectively. Here we directly employ the number of points in each voxel to represent its point density as all voxels have the same size. The formula can be explained as follows. Whenever V_m^i and/or V_n^i are empty (i.e., $N_m^i = 0$ or $N_n^i = 0$), $s(V_m^i, V_n^i) = 0$. Otherwise, when two corresponding voxels contain similar local shape, their centroids should be close to each other, making $s(V_m^i, V_n^i)$ large. When N_m^i and/or N_n^i are large, $s(V_m^i, V_n^i)$ is large also as the estimated centroid(s) are more accurate. By this, the formula encourages to find matches based on denser parts of input scans when the scans are irregularly sampled.

The overall similarity score between D_m and D_n can be obtained by accumulating the similarity value for every pair of corresponding voxels:

$$S(D_m, D_n) = \sum_{i=1}^{K \times K \times K} s(V_m^i, V_n^i) \quad (2)$$

3.4 SGC Generation Parameters

The SGC descriptor has two generation parameters: (i) the support radius R ; and (ii) the voxel grid resolution K . According to our experiments, we choose $R = 20 pr$ as a tradeoff between the descriptiveness and sensitivity to occlusion, where pr denotes the point cloud resolution (i.e., average shortest distance among neighboring points in the scan). And we choose $K = 8$ as a tradeoff between the descriptiveness and efficiency since a larger K increases the descriptiveness and computational cost simultaneously. Note that in these experiments, we let the LRF and the descriptor have the same support radius, i.e., $r = R$.

4 Partial Shape Matching using SGC

In this section, we describe the general pipeline to match two scans using SGC descriptors and propose a descriptor saliency measure for improving shape matching performance. We also highlight the advantage of using SGC descriptors for matching supports that are close to scan boundary.

4.1 General Shape Matching Pipeline

Given a data scan S_d and a reference scan S_r , the goal of shape matching between S_d and S_r is to find a rigid transform on S_d to align it with S_r . By employing the SGC descriptors, we can find such a transform with following steps:

1) *Represent Scans with SGC Descriptors.* We first conduct a uniform sampling on each of S_d and S_r to generate M feature points that cover the whole scan surface. Next, for each feature point p , we construct the LRF and SGC descriptor for the support around p . By this, we represent each of S_d and S_r with M descriptor vectors and the corresponding LRFs, see Figure 2(a&b).

2) *Generate Transform Candidates.* When a point on S_d corresponds to another point on S_r , their associated SGC descriptors should be similar to each other. Hence, we compare each feature descriptor of S_d with each feature descriptor of S_r by calculating a similarity score using Eq. 2. A feature point on S_d and its closest feature point on S_r are considered as a match if the similarity score is higher than a threshold. Each match generates a rigid transform candidate (i.e., a 4×4 transformation matrix) by aligning the associated LRFs.

3) *Select the Optimal Transform.* By matching the descriptors of S_d and S_r , we obtain a number of candidate transforms. We sort these transforms based on the descriptor similarity score and then pick the top five candidates with the highest scores. We apply each of the five selected transforms on S_d to align it with S_r . We evaluate the transform by computing a scan overlap ratio. We first find all point-to-point correspondences by checking if the distance between a point on transformed S_d and a point on S_r is sufficiently small, and further compute the overlap ratio as the number of corresponding points divided by the total number of points in S_d or S_r (smaller one). We select the transform that ensures the largest overlap ratio as the optimal one, see Figure 2(c&d).

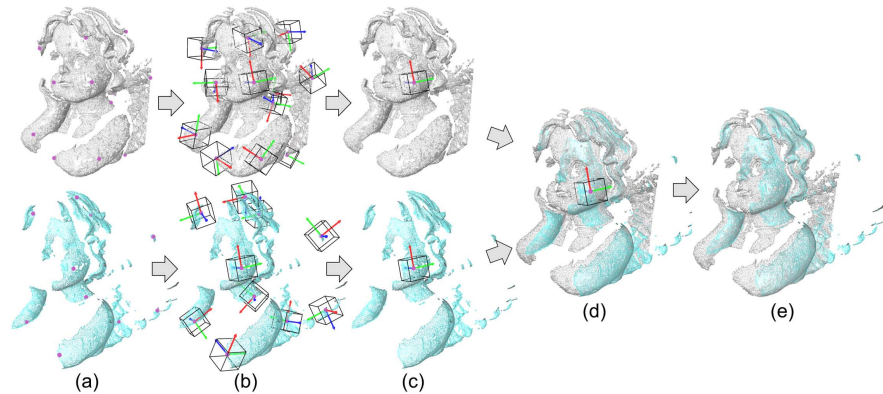


Fig. 2. Matching two scans using SGC descriptors: (a) sampled feature points (in purple) on two input scans (only part of samples are shown for clarity); (b) calculated LRFs and descriptors; (c) a pair of matched descriptors; (d) match the two scans based on aligning the associated LRFs; and (e) refine the scan alignment using ICP.

4) *Refine the Scan Alignment.* Optionally, we can apply iterative closest point (ICP) to refine the alignment generated by the selected optimal transform, see Figure 2(e). By comparing Figure 2(d&e), we can see that the transform calculated by aligning LRFs is very close to the one refined using ICP.

4.2 Improve Shape Matching using Descriptor Saliency

To ensure corresponding points to be found on different scans, we need to sample a large number of feature points on each scan, e.g., $M = 1000$ in our experiments. However, among the M descriptors on a single scan, there could exist some descriptors close to one another since their corresponding supports are similar, see Figure 3(a). Moreover, among descriptors from all input scans, there could exist a larger number of descriptors with high similarities, see Figure 3(a-c).

Our observation is that when there exist a large number of descriptors with high similarities, it means their corresponding supports are less distinctive (e.g., flat or spherical shape), see the zooming views in Figure 3(a). Thus, it has a lower chance to match the scans correctly by using such supports and their descriptors. On the other hand, when a descriptor is quite different from others, it means its support is distinctive (see the top zooming views in Figure 3(b&c)).

Inspired by this observation, we propose a measure of *descriptor saliency* to improve the shape matching performance and compute it based on a descriptor-graph. The key idea is to find descriptors (and the corresponding supports) that are distinctive by measuring their saliency and apply these descriptors to find corresponding feature points. We first describe our approach to build a descriptor-graph, present our definition on the descriptor saliency, and then show how we apply the descriptor saliency to enhance shape matching.

Build a Descriptor-Graph. For a given reference scan S_r , we build a descriptor-graph for all the descriptors sampled from S_r based on their similarities computed using Eq. 2. Formally, let $G = (V, E)$ be a descriptor-graph, each node $u \in V$ represents an SGC descriptor on S_r , while each directed edge $(u, v) \in E$ represents that v is one of k -nearest neighbours (k -NN) of u in the descriptor similarity space. Note that we do not require u also to be one of k -NN of v , which means there may not exist a directed edge (v, u) in G .




Fig. 3. Supports on three different scans of a CHEF model, where feature points are rendered in pink. The correspondence between a scan support on the left and its zooming view on the right is indicated by the same 2D box color.

To build such a graph, a straightforward way is to exhaustive search all descriptors on S_r to retrieve k-NN for each descriptor in G . However, this approach is time-consuming, especially when G is large. We speed up the creation of the graph following [28], and the basic idea is to initially fill the nearest neighbors by randomly sampling descriptors in G , and iteratively optimize the nearest neighbors locally via similarity propagation and random search until convergence.

Define Descriptor Saliency. We define descriptor saliency as the distinctiveness among a set of given descriptors. The larger difference between a descriptor and others, the higher its saliency. Thus, we measure saliency of a descriptor D_i in a descriptor-graph G using $sali(D_i) = \frac{1}{1+e^{(I_i-\bar{I})}}$, where I_i denotes the number of nodes in G that considers D_i as a k-NN and \bar{I} is the mean value of all I_i that is larger than zero. Note that although D_i has k nearest neighbors in G , these neighbors could be very different from D_i . By fixing k , the value I_i can reveal how many descriptors are close to D_i (i.e., D_i 's distinctiveness). Figure 4 shows descriptor saliency in a simple descriptor-graph with $k = 3$.

Shape Matching with Descriptor Saliency. For a given reference scan S_r , we first create a descriptor-graph G_r for it and compute a saliency value for every descriptor D_r^i in G_r using $sali(D_r^i)$. For a given descriptor on the data scan S_d , say D_d^j , we enhance the similarity score between D_d^j and D_r^i by using $sali(D_r^i)$, i.e., $\bar{S}(D_d^j, D_r^i) = sali(D_r^i)^\alpha S(D_d^j, D_r^i)$, where α is a weight to control the impact of saliency on the descriptor similarity. We set $\alpha = 0.2$ in our experiments.

Intuitively, we can find the descriptor on S_r corresponding to D_d^j on S_d by simply comparing every D_r^i on S_r with D_d^j and selecting the one with the largest $\bar{S}(D_d^j, D_r^i)$. We speed up the search of the corresponding descriptor by taking advantage of G_r with the idea of leveraging existing matches to find better ones. This is achieved by randomly selecting a set of nodes in G_r and updating the nodes by a few iterations of similarity propagation and random search [29], guided by the similarity score (using Eq. 2) between D_d^j and the nodes. After obtaining a small set of descriptors on S_r that are similar to D_d^j , we conduct re-ranking using $\bar{S}(D_d^j, D_r^i)$ to select the final correspondence.

We have illustrated applying descriptor saliency for shape matching between a pair of scans. Descriptor saliency is more suitable for shape matching among a number of scans, with following changes. First, we build a large descriptor-graph G for descriptors from all the scans. Second, we compare a descriptor on scan S_m




Fig. 4. An example descriptor-graph (outdegree = 3 for every node).

with nodes in G that are not from S_m . By this, the larger the number of scans, the higher shape matching performance can be improved by descriptor saliency.

4.3 Matching Supports Close to Scan Boundary

Depth scans captured from a certain view are mostly incomplete due to a limited viewing angle, sensor noise, and occlusion. This results in a surface boundary for a scan. Matching supports close to the boundary is a challenging task. First, the support is likely to be incomplete, see examples in Figure 2(b). This affects an LRF’s repeatability since support is the only input to construct the LRF. Further, deviation of the LRF affects the construction of the descriptor since support partitioning is performed within the LRF. Second, the incomplete support directly affects the construction of the descriptor since voxels locating at the missing part(s) become empty, where no shape feature can be extracted.

Due to the above challenges, many existing descriptors are sensitive to the boundary points according to the evaluation in [17]. Therefore, boundary points are usually ignored when applying existing descriptors to partial shape matching [30, 7], assuming that there is sufficient non-boundary scan surface for the matching. On the other hand, matching boundary points will improve the chance to correctly align different scans, especially when the scan overlap is small.

Our SGC descriptor is especially suitable for handling boundary points for shape matching. First, the centroid feature that SGC employs is robust against noise and varying point density, which usually happen at scan boundary. Second, our descriptor comparison scheme allows matching descriptors computed from either a complete or an incomplete support, see Figure 5. Third, we allow using two different radii for constructing the LRF and the descriptor, i.e., $r \leq R$, see supports with varying sizes in Figure 5(left). By this, a smaller yet complete support can be employed for constructing a repeatable LRF while a larger support allows encoding more (complete or incomplete) local shape for constructing the descriptor. Based on our experiments, we find that $r = 0.5R$ achieves the best performance for matching boundary points when setting $R = 20 pr$.




Fig. 5. (left) Match a support containing holes (in gray scan) with a support close to boundary (in cyan scan) using SGC descriptors; (right) aligned scans and supports.

5 Performance of the SGC Descriptor

This section evaluates the robustness of SGC with respect to various nuisances, including noise, varying point density, distance to scan boundary, and occlusion. We compare SGC with three state-of-the-art descriptors that work on point cloud data: spin image (SI) [24], 3DSC [10] and SHOT [11]. Table 1 presents a detailed description of the parameter settings.

Table 1. Parameter settings of the four descriptors.

	Radius	Bin feature	Dimensionality	Length
SGC	20pr	Geometric centroid	$8 \times 8 \times 8 \times 2$	1024
SHOT	20pr	Histogram of normals	$8 \times 2 \times 2 \times 10$	320
3DSC	20pr	Point density	$15 \times 11 \times 12$	1980
SI	20pr	Point density	15×15	225

We perform the experiments on three publicly available datasets: the Bologna dataset [31], UWA dataset [30], and Queen’s dataset [32]. Unlike the Bologna dataset that synthesizes complete object models to generate scenes, the scenes in the UWA and Queen’s dataset contain partial shape of object models. We employ the Bologna dataset to evaluate the descriptors’ performance with respect to noise and varying point density (Subsection 5.1 & 5.2), the UWA dataset to evaluate the descriptors’ performance with respect to distance to scan boundary and occlusion (Subsection 5.3 & 5.4), and the Queen’s dataset to evaluate improved performance by using descriptor saliency (Subsection 5.5).

We compare the descriptors’ performance using RP curves [33]. In detail, we randomly select 1000 feature points in each model and find their corresponding points in the scenes via the physical nearest neighbouring search. By matching the scene features against the model features using each of the four descriptors, an RP curve of the descriptor is generated.

5.1 Robustness to Noise

To evaluate robustness of the descriptors against noise, we add four different levels of Gaussian noise with standard deviations of 0.1, 0.3, 0.5, and 1.0 pr to each scene. The RP curves of the four descriptors are presented in Figure 6(a-d). Thanks to the robust centroid feature, the RP curves show that SGC performs the best under all levels of noise, followed by SHOT and 3DSC.

5.2 Robustness to Varying Point Density

To evaluate robustness of the descriptors with respect to varying point density, we downsample the noise free scenes to 1/2, 1/4 and 1/8 of their original point density (pd). The RP curves in Figure 6(e-g) show that SGC outperforms all other descriptors under all levels of downsampling. Figure 6(h) shows that SGC performs the best when the input scans are downsampled and contain noise.

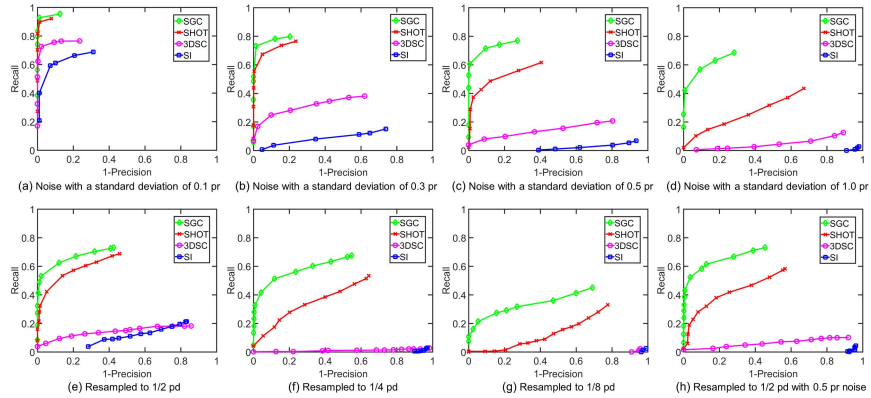


Fig. 6. RP curves of the four descriptors in the presence of (a-d) noise, (e-g) point cloud downsampling, and (h) their combination.

5.3 Robustness to Distance to Scan Boundary

We perform experiments for feature points within different ranges of distance to the boundary, i.e., $(0, 0.25R]$, $(0.25R, 0.5R]$, $(0.5R, 0.75R]$, and $(0.75R, R]$. Note that we set tuned $r = 0.5R$ for SGC and $r = R$ for all the other descriptors. Thanks to the varying support radius and descriptor comparison scheme, Figure 7 shows that SGC achieves the best performance for all the four cases.




Fig. 7. RP curves of feature points in different ranges of distance to the scan boundary.

5.4 Robustness to Occlusion

To evaluate performance of the descriptors under occlusion, we group sampled feature points into two categories following [17], i.e., $(60\%, 70\%]$ and $(70\%, 80\%]$ occlusions. Figure 8(a&b) shows that SGC outperforms all the other descriptors with a large margin since SGC allows handling feature points at scan boundary.

5.5 Effectiveness of Descriptor Saliency

To demonstrate effectiveness of descriptor saliency, we compare our shape matching approach with an exhaustive search to find corresponding feature points.




Fig. 8. (a&b) RP curves about occlusion. (c) CMC curves about descriptor saliency.

First, we build a descriptor-graph for descriptors sampled from all the five models in the Queen’s dataset [32] with $k = 16$. Next, we randomly select 1000 feature points on a scene and calculate their SGC descriptors. For each scene descriptor, we retrieve its neighbours by searching the descriptor-graph with saliency or exhaustive searching all the model descriptors. Here, we concern how many neighbours we need to retrieve to ensure the corresponding descriptor is included. Figure 8(c) shows standard Cumulated Matching Characteristics (CMC) curves [34] by using the two approaches. The curves show that descriptor saliency brings a certain amount of improvement in shape matching. In addition, descriptor-graph speeds up the search of corresponding descriptors, where each query process takes $0.5ms$, much faster than the exhaustive search ($62ms$).

6 Applications

3D Object/Scene Reconstruction. To reconstruct a more complete model from a set of scans, we build a descriptor-graph for all the scans. As the graph has encoded k-NN for each descriptor (and the feature point), we search the corresponding feature point (and its associated scan ID) locally within the k-NN, and align the two scans based on the correspondence and merge them into a larger point cloud. We keep aligning each of the remaining scans with the point cloud and merging them until all scans are registered. Figure 9 shows two objects and one scene reconstructed by our approach on different datasets [11, 35].

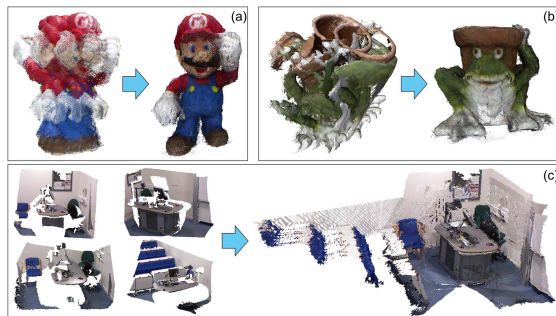
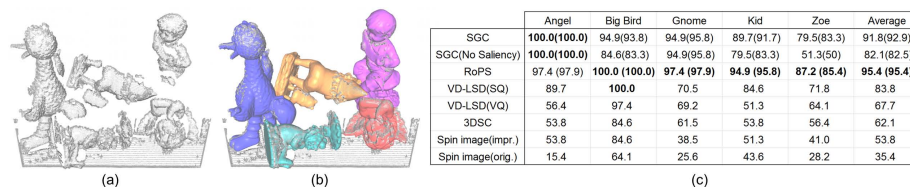


Fig. 9. Our reconstruction results. (a) SUPER MARIO; (b) FROG; and (c) STAGE SCENE.

3D Object Recognition. We conduct this experiment on the challenging Queen’s dataset [32]. To represent the model library well with SGC, we remove the noise in each model point cloud and build a descriptor-graph for descriptors sampled from all the models. For a give scene scan, we also sample a number of SGC descriptors. By searching a corresponding descriptor in the graph for a given scene descriptor, we know the correspondence between a model in the library and a partial scene, thus recognizing the object in the scene scan. Note that we recognize a single object at a time and segment the object once recognized.

Figure 10(a&b) show the recognition result on an example scene. Figure 10(c) shows that SGC based algorithm outperforms most existing methods including VD-LSD [32], 3DSC [10] and spin image [24] based algorithms. RoPS based algorithm is the current best 3D object recognition approach and it achieves slighter better performance than SGC with additional mesh information of the scene scans. In particular, the performance of our algorithm without using descriptor saliency decreases about 10%, indicating the usefulness of the saliency.



	Angel	Big Bird	Gnome	Kid	Zoe	Average
SGC	100.0(100.0)	94.9(93.8)	94.9(95.8)	89.7(91.7)	79.5(83.3)	91.8(92.9)
SGC(No Saliency)	100.0(100.0)	84.6(83.3)	94.9(95.8)	79.5(83.3)	51.3(50)	82.1(82.5)
RoPS	97.4 (97.9)	100.0 (100.0)	97.4 (97.9)	94.9 (95.8)	87.2 (85.4)	95.4 (95.4)
VD-LSD(SQ)	89.7	100.0	70.5	84.6	71.8	83.8
VD-LSD(VQ)	56.4	97.4	69.2	51.3	64.1	67.7
3DSC	53.8	84.6	61.5	53.8	56.4	62.1
Spin Image(impr.)	53.8	84.6	38.5	51.3	41.0	53.8
Spin Image(orig.)	15.4	64.1	25.6	43.6	28.2	35.4

Fig. 10. Recognition results on the Queen’s dataset. (a) An example scene; (b) our recognition results; and (c) recognition rates of the five models (values in brackets are the results on the whole dataset while others are the results on the subset as in [32]).

7 Conclusion

We have presented a novel SGC descriptor for matching partial shapes represented by 3D point clouds. SGC integrates three novel components: 1) a local shape description that encodes robust geometric centroid features; 2) a descriptor comparison scheme that allows comparing supports with missing parts; and 3) a descriptor saliency measure that can identify distinct descriptors. By this, SGC is robust against various nuisances in point cloud data when performing partial shape matching. We have demonstrated SGC’s performance by comparisons with state-of-the-art descriptors and two partial matching applications.

Acknowledgments

This work is supported in part by the Fundamental Research Funds for the Central Universities (WK0110000044), Anhui Provincial Natural Science Foundation (1508085QF122), National Natural Science Foundation of China (61403357, 61175057), and Microsoft Research Asia Collaborative Research Program.

References

1. Aiger, D., Mitra, N.J., Cohen-or, D.: 4-points congruent sets for robust pairwise surface registration. *ACM Trans. on Graphics (Proc. of SIGGRAPH)* **27** (2008) Article 85.
2. Bariya, P., Nishino, K.: Scale-hierarchical 3D object recognition in cluttered scenes. In: *CVPR*. (2010) 1657–1664
3. Donoser, M., Riemenschneider, H., Bischof, H.: Efficient partial shape matching of outer contours. In: *ACCV*. (2009) 281–292
4. Rodolà, E., Cosmo, L., Bronstein, M.M., Torsello, A., Cremers, D.: Partial functional correspondence. *Computer Graphics Forum* (2016)
5. Mian, A.S., Bennamoun, M., Owens, R.A.: A novel representation and feature matching algorithm for automatic pairwise registration of range images. *International Journal of Computer Vision* **66** (2006) 19–40
6. Wu, H.Y., Zha, H., Luo, T., Wang, X.L., Ma, S.: Global and local isometry-invariant descriptor for 3D shape comparison and partial matching. In: *CVPR*. (2010) 438–445
7. Guo, Y., Sohel, F., Bennamoun, M., Lu, M., Wan, J.: Rotational projection statistics for 3D local surface description and object recognition. *International Journal of Computer Vision* **105** (2013) 63–86
8. Song, P., Chen, X.: Pairwise surface registration using local voxelizer. In: *Pacific Graphics*. (2015) 1–6
9. Johnson, A.E.: *Spin-Images: A Representation for 3-D Surface Matching*. PhD thesis, Robotics Institute, Carnegie Mellon University, Pittsburgh, PA (1997)
10. Frome, A., Huber, D., Kolluri, R., Bülow, T., Malik, J.: Recognizing objects in range data using regional point descriptors. In: *ECCV*. (2004) 224–237
11. Tombari, F., Salti, S., Stefano, L.D.: Unique signatures of histograms for local surface description. In: *ECCV*. (2010) 356–369
12. Guo, Y., Sohel, F., Bennamoun, M., Wan, J., Lu, M.: A novel local surface feature for 3D object recognition under clutter and occlusion. *Information Sciences* **293** (2015) 196–213
13. Guo, Y., Bennamoun, M., Sohel, F., Lu, M., Wan, J.: 3D object recognition in cluttered scenes with local surface features: A survey. *IEEE Trans. on Pattern Analysis and Machine Intelligence* **36** (2014) 2270–2287
14. Hilaga, M., Shinagawa, Y., Kohmura, T., Kunii, T.L.: Topology matching for fully automatic similarity estimation of 3D shapes. In: *SIGGRAPH*. (2001) 203–212
15. Chao, M.W., Lin, C.H., Chang, C.C., Lee, T.Y.: A graph-based shape matching scheme for 3D articulated objects. *Computer Animation and Virtual Worlds* **22** (2011) 295–305
16. Osada, R., Funkhouser, T., Chazelle, B., Dobkin, D.: Shape distributions. *ACM Trans. on Graphics* **21** (2002) 807–832
17. Guo, Y., Bennamoun, M., Sohel, F., Lu, M., Wan, J., Kwok, N.M.: A comprehensive performance evaluation of 3D local feature descriptors. *International Journal of Computer Vision* **116** (2016) 66–89
18. Gal, R., Cohen-Or, D.: Salient geometric features for partial shape matching and similarity. *ACM Trans. on Graphics* **25** (2006) 130–150
19. Albarelli, A., Rodolà, E., Torsello, A.: Loosely distinctive features for robust surface alignment. In: *ECCV*. (2010) 519–532
20. Petrelli, A., Stefano, L.D.: On the repeatability of the local reference frame for partial shape matching. In: *ICCV*. (2011) 2244–2251

21. Hetzel, G., Leibe, B., Levi, P., Schiele, B.: 3D object recognition from range images using local feature histograms. In: CVPR. Volume 2. (2001) 394–399
22. Yamany, S.M., Farag, A.A.: Surface signatures: An orientation independent free-form surface representation scheme for the purpose of objects registration and matching. *IEEE Trans. on Pattern Analysis and Machine Intelligence* **24** (2002) 1105–1120
23. Kokkinos, I., Bronstein, M.M., Litman, R., Bronstein, A.M.: Intrinsic shape context descriptors for deformable shapes. In: CVPR. (2012) 159–166
24. Johnson, A.E., Hebert, M.: Using spin images for efficient object recognition in cluttered 3D scenes. *IEEE Trans. on Pattern Analysis and Machine Intelligence* **21** (1999) 433–449
25. Zhong, Y.: Intrinsic Shape Signatures: A shape descriptor for 3D object recognition. In: 12th International Conference on Computer Vision Workshops. (2009) 689–696
26. Song, P.: Local Voxelizer: A shape descriptor for surface registration. *Computational Visual Media* **1** (2015) 279–289
27. Pottmann, H., Wallner, J., Huang, Q.X., Yang, Y.L.: Integral invariants for robust geometry processing. *Computer Aided Geometric Design* **26** (2009) 37–60
28. Barnes, C., Shechtman, E., Goldman, D.B., Finkelstein, A.: The generalized patch-match correspondence algorithm. In: ECCV. (2010) 29–43
29. Gould, S., Zhao, J., He, X., Zhang, Y.: Superpixel graph label transfer with learned distance metric. In: ECCV. (2014) 632–647
30. Mian, A.S., Bennamoun, M., Owens, R.: Three-dimensional model-based object recognition and segmentation in cluttered scenes. *IEEE Trans. on Pattern Analysis and Machine Intelligence* **28** (2006) 1584–1601
31. Tombari, F., Salti, S., Stefano, L.D.: Unique shape context for 3D data description. In: ACM Workshop on 3D Object Retrieval. (2010) 57–62
32. Taati, B., Greenspan, M.: Local shape descriptor selection for object recognition in range data. *Computer Vision and Image Understanding* **115** (2011) 681–694
33. Mikolajczyk, K., Schmid, C.: A performance evaluation of local descriptors. *IEEE Trans. on Pattern Analysis and Machine Intelligence* **27** (2005) 1615–1630
34. Wang, X., Doretto, G., Sebastian, T., Rittscher, J., Tu, P.: Shape and appearance context modeling. In: ICCV. (2007) 1–8
35. Mellado, N., Aiger, D., Mitra, N.J.: Super 4PCS fast global pointcloud registration via smart indexing. *Computer Graphics Forum (Proc. of SGP)* **33** (2014) 205–215