

A Curvature-based Approach for Multi-scale Feature Extraction from 3D Meshes and Unstructured Point Clouds

Huy Tho Ho and Danny Gibbins

Sensor Signal Processing Group

School of Electrical and Electronic Engineering

The University of Adelaide, Adelaide, Australia

{huyho, danny}@eleceng.adelaide.edu.au

Abstract

A framework for extracting salient local features from 3D models is presented in this paper. In the proposed method, the amount of curvature at a surface point is specified by a positive quantitative measure known as the curvedness. This value is invariant to rigid body transformation such translation and rotation. The curvedness at a surface position is calculated at multiple scales by fitting a manifold to the local neighbourhoods of different sizes. Points corresponding to local maxima and minima of curvedness are selected as suitable features and a confidence measure of each keypoint is also calculated based on the deviation of its curvedness from the neighbouring values. The advantage of this framework is its applicability to both 3D meshes and unstructured point clouds. Experimental results on a different number of models are shown to demonstrate the effectiveness and robustness of our approach.

I. INTRODUCTION

As surface acquisition methods such as LADAR or range scanners are becoming more popular, there is an increasing interest in the use of three-dimensional geometric data in various computer vision applications. However, the processing of 3D datasets such as range images is a demanding job due to not only the huge amount of surface information but also the noise and non-uniform sampling introduced by the sensors or the reconstruction process. It is therefore desirable to have a more compact intermediate representation of 3D images that can be used efficiently in computer vision tasks such as 3D shape modelling, surface registration or object recognition.

One of the most popular approaches is based on using local descriptors, or signatures that describe local surface regions, to represent a surface [1]–[4]. This strategy has been proven to be robust to partial occlusion, clutter and intra-class variation [5]. However, in those techniques, the surface locations used for estimating local descriptors are either selected exhaustively at each point or randomly from the data [1], [2]. In the case of exhaustive selection, it is very inefficient because of the redundancy in areas with little shape variation. In the case of randomly selection, distinctive geometric structures may be missed thus




Fig. 1. Local features detected on the ‘Dragon’ model. The colormap of the keypoints is from blue to red corresponding to low to high confidence values. Figures are best seen in color.

reducing the accuracy of the algorithm [5]. Therefore, it is very important to be able to have a principled way to sample a representative set of feature points from the surface [6].

This paper addresses the above issue by proposing a multi-scale feature extraction algorithm using a rotation and translation invariant local surface curvature measure known as the curvedness. It is a positive number that captures the amount of curvature in a local region around a surface point [7]. Different values of the curvedness of a point are calculated at multiple scales by fitting a surface to its neighbourhood of different sizes. A set of reliable salient feature points is formed by finding the set of extrema from the scale-space representation of the input surface. We also introduce a method for evaluating the confidence of each keypoint based on the deviation of its curvedness from the neighbouring values. The proposed approach is tested on a variety of 3D models with different noise levels in order to show the high repeatability of the detected features. Finally, registration results obtained by using our feature selection framework with spin-images [2] as local descriptors are also included to demonstrate the robustness and effectiveness of the approach. Figure 1 shows the features extracted by our algorithm from the Stanford ‘Dragon’ model [8].

The remainder of the paper is organized as follows. Section II describes related work in the area. Surface curvature and method of estimation are described in Section III. In Section IV, we present in detail how to perform the proposed multi-scale feature extraction algorithm. Our experimental results are presented in Section V. The conclusions are finally given in Section VI.

II. RELATED WORK

A. Scale-space Representation

In 2D image domain, multi-scale feature extraction is a well-established problem [9]–[11]. The fundamental idea of the scale-space representation, first introduced in 1983 [9], is to transform an input signal $f(\mathbf{x}) : \mathbf{R}^d \rightarrow \mathbf{R}$ at different scales t as $L(\mathbf{x}, t) : \mathbf{R}^d \times \mathbf{R}_+ \rightarrow \mathbf{R}$. The Gaussian scale-space representation can be obtained by convolving $f(\mathbf{x})$ with Gaussian kernels G of increasing width t [10]

$$L(\mathbf{x}, t) = G(\mathbf{x}, t) \otimes f(\mathbf{x}) \quad (1)$$

Figure 2 shows an example of applying Gaussian kernels of increasing widths to a 2D image in order to create its scale-space representation.



Fig. 2. Scale-space representation of a 2D image with increasing Gaussian kernel width.

The connection between the Gaussian scale-space and the diffusion equation was obtained in [12]

$$\frac{\partial L(\mathbf{x}, t)}{\partial t} = \Delta L(\mathbf{x}, t) \equiv \sum_i^d \frac{\partial^2 L(\mathbf{x}, t)}{\partial x_i^2} \quad (2)$$

where Δ denotes the Laplacian operator. There are many methods that extends the above idea of 2D scale-space representation to point-sampled surfaces. One method is the iterative Laplacian smoothing [13] that used the graph Laplacian L_g to replace the continuous Laplacian operator Δ

$$L_g f(\mathbf{x}_i) = \sum_{j, \{i,j\} \in \mathbf{E}} (f(\mathbf{x}_i) - f(\mathbf{x}_j)) w_{ij} \quad (3)$$

where \mathbf{x}_i is a set of vertices of a graph, the summation is over graph edges (i, j) in the edge set \mathbf{E} and w_{ij} are positive edge weights. However, the drawback of this method is that its smoothing kernel often produces surface deformation artifacts such as volume shrinkage and thus may incorrectly change the surface geometry [5], [14].

In [14], a surface variation measure, $\sigma_n(x)$, was proposed

$$\sigma_n(\mathbf{x}) = \frac{\lambda_0}{\lambda_0 + \lambda_1 + \lambda_2} \quad (4)$$

where λ_i 's are eigenvalues of the sample covariance matrix evaluated in a local n -point neighbourhood of the sample point \mathbf{x} . The scale of an extracted feature is chosen to be the neighbourhood size for which the corresponding σ_n gets a local extremum. However, σ_n is very sensitive to noise and this approach requires heuristic pre-smoothing procedures to be applied to the surface [5].

A recent work proposed by Novatnack and Nishino [15] aims at detecting multi-scale corner and edge features from 3D meshes. In this approach, a 2D representation of the original surface called the *normal map* is created by interpolating over the surface normals at each 2D-embedded vertex. Corner and edge detectors are derived using the first and second-order partial derivatives of the normal map in the horizontal (s) and vertical direction (t). The disadvantage of this approach is that the 2D normal map can only be created from a 3D mesh with connectivity information. It also suffers from another limitation that good surface normals must be available in order to construct the scale-space representation.

Recently, Flint *et al.* [16] proposed the use of a 3D version of the Hessian to measure the distinctiveness of candidate interest points. The disadvantage of this method is the cost of resampling regularly in space throughout the data. Another drawback is the stability of the Hessian determinant used for searching local maxima in noisy data.

Surface curvature has been used extensively in the literature for mesh simplification and smoothing [17], object recognition [18]–[20] and mesh segmentation [21]. However, there is a lack of a systematic approach in extracting salient local features




Fig. 3. The (α, β) coordinate of the surface point q relative to p .




Fig. 4. A spin-image extracted at a vertice on the face of the ‘Chef’ model. The location of the vertice is marked in red. Surface points contributing to the spin-image are highlighted by blue line segments oriented along the surface normals at these positions.

or keypoints from an input surface using its local curvature information at multiple scales. In this paper, we address the above issue by proposing a multi-scale curvedness-based approach.

B. Spin-Images

In order for the paper to be self-contained, the spin-image algorithm, proposed by Johnson *et al.* [2], is briefly discussed in this section. Each spin-image is a local surface descriptor calculated at an oriented point (\mathbf{p}, \mathbf{n}) (3D point with normal vector) by encoding two of the three cylindrical coordinates of the its surrounding points (Figure 3). The spin-image X for a surface point \mathbf{p} is a 2D histogram in which each pixel is a bin that stores the number of neighbours that are a distance α from \mathbf{n} and a depth β from its tangent plane \mathbf{P} . Figure 4 shows the ‘Chef’ model [22] and a spin-image extracted at a vertice on its face.

The similarity of two spin-images is measured by calculating their correlation coefficient [2]. The registration of two different views are performed by finding the correspondences between points on the surfaces using the similarity of their spin-images. Correspondence pairs with highest similarity that are geometrically consistent are selected to estimate the rigid transformation that registers the surfaces. As the surface positions chosen to compute spin-images are selected randomly, this may reduce the accuracy of the registration as a result of missing important geometric structures. By incorporating our feature selection




Fig. 5. Principal curvatures at a point \mathbf{P} on a surface.

framework into the spin-image algorithm, we can improve not only its robustness by using salient features in the matching process but also the speed as the number of features is significantly smaller than the number of randomly selected points.

III. LOCAL SURFACE CURVATURE

A. Principal Curvatures

If we let \mathbf{u}_p be a tangent vector to a regular surface $M \subset \mathbf{R}^3$ at point $\mathbf{p} \in M$ with $\|\mathbf{u}_p\| = 1$, the normal curvature of M in the direction of \mathbf{u}_p is defined as

$$\mathbf{k}(\mathbf{u}_p) = S(\mathbf{u}_p) \cdot \mathbf{u}_p \quad (5)$$

where $S(\mathbf{u}_p) = -D_{\mathbf{u}_p} \mathbf{U}$ is the shape operator. \mathbf{U} is the field of surface normal vectors defined in a neighbourhood of \mathbf{p} and $D_{\mathbf{u}_p} \mathbf{U}$ is the derivative of \mathbf{U} with respect to \mathbf{u}_p [23].

The maximum and minimum values, k_1 and k_2 respectively, of the normal curvature $\mathbf{k}(\mathbf{u}_p)$ are called principal curvatures of M at point \mathbf{p} . These values measure the maximum and minimum bending of M at \mathbf{p} [7]. The principal curvatures are related to the two classic shape measures, the Gaussian curvature K and mean curvature H , by

$$k_1 = H + \sqrt{H^2 - K} \quad (6)$$

$$k_2 = H - \sqrt{H^2 - K} \quad (7)$$

Figure 5 shows the principal curvatures for a point \mathbf{p} with the normal vector \mathbf{n} on a surface. The curvature is positive if the curve turns in the same direction as the surface normal \mathbf{n} . Otherwise, it will be negative. Figure 6 shows the maximum and minimum principal curvatures estimated for the Stanford ‘Buddha’ model. The colormap is shown in Figure 6c.

B. Curvedness

In our feature extraction framework, the geometric aspect of a 3D model is defined using a bending energy measure of the surface called the curvedness. The curvedness at a point \mathbf{p} on a surface can be estimated as [7], [18]

$$c_p = \sqrt{\frac{k_1^2 + k_2^2}{2}} \quad (8)$$

The curvedness can be used to indicate how highly or gently curved a surface is [21]. A single curvedness value is sufficient to describe the local surface curvature, whereas both the Gaussian and mean curvatures are necessary for the same task [18]. The curvedness is zero only for planar patches as $k_1 = k_2 = 0$ in this case. Figures 7a and 7b plot the curvedness of two 3D models, ‘Buddha’ and ‘Dragon’, respectively. The color map is shown in Figure 7c.




Fig. 6. Maximum and minimum principal curvatures estimated for the ‘Buddha’ model. The colormap is from blue to red corresponding to low to high values of the curvature. Figures are best seen in color.

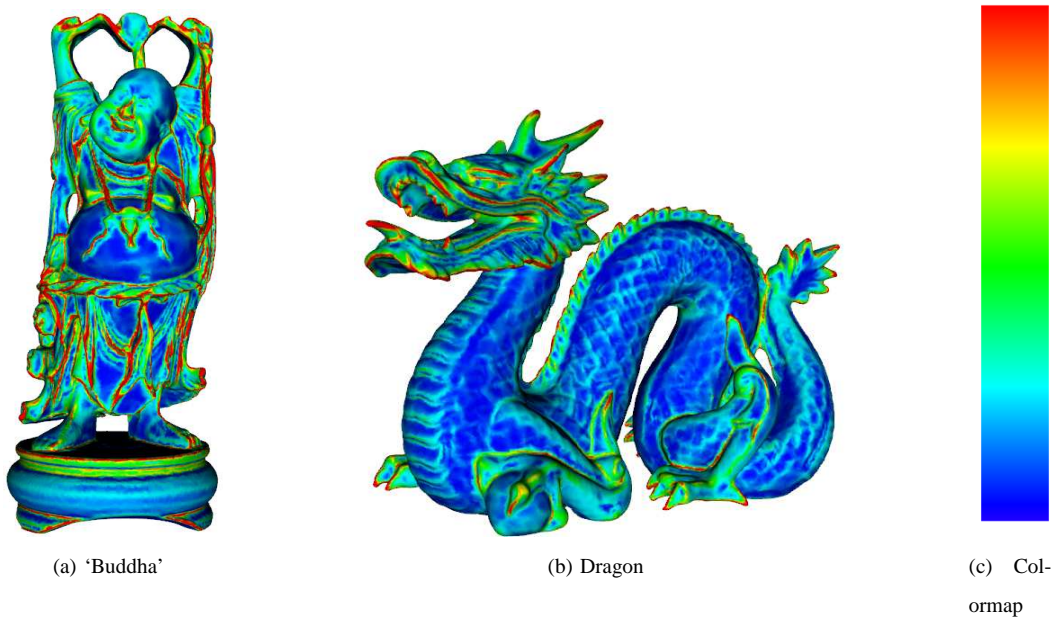


Fig. 7. Curvedness of the ‘Buddha’ and ‘Dragon’ models. The colormap is from blue to red corresponding to low to high values of the curvedness.

IV. MULTI-SCALE FEATURE EXTRACTION

A. Feature Point Selection at Multi-scale

As we have already discussed in Section II, there are equivalent 3D versions of the 2D Gaussian smoothing kernels in order to estimate the 3D scale space representation of a surface. However, these methods usually produce undesired artifacts such as changes in the geometric structures of the models. In our framework, the scale of a point on the surface is defined as the size of the neighbourhood that we use to collect points as the input to the fitting. For unstructured point clouds, the scale r can be chosen as either Euclidean distance or geodesic distance [23]. All the surface points that are closer to the fitting point \mathbf{p}




Fig. 8. An example of rings. The fitting vertex is marked with blue color.




Fig. 9. Curvedness of the ‘Buddha’ model estimated at different scale levels r_k . It can be seen that increasing the scale is similar to applying a smoothing filter to the 3D surface before calculating the surface curvature. The colormap is from blue to red corresponding to low to high values of the curvedness.

than this distance will be picked. In this paper, we use the Euclidean distance to specify the radius of the neighbourhood. For 3D meshes, the scale r is chosen as the number of *rings* surrounding the fitting vertex \mathbf{v} . The scale r can also be called the radius of the neighbourhood. The first ring contains all the direct neighbours of \mathbf{v} that are vertices in the opposite of the edges started from \mathbf{v} . The second ring contains all the direct neighbours of the vertices in the first ring and so on. In the scope of this paper, we only use meshes as inputs to our feature selection algorithm and the number of rings is used as the scale level. Figure 8 shows an example of a vertex and its neighbourhoods of radius $r = 1$ and $r = 2$.

Our multi-scale curvedness-based feature selection algorithm is outlined in Algorithm 1. It can be seen that increasing the size of the local neighbourhood is similar to applying a smoothing filter but it avoids making direct changes to the 3D surfaces [14]. Thus, in addition to the ability to detect features at multiple scales, another benefit of this approach is to reduce the effect of noise on the models. It is also worth noting that the scales r_k where keypoints selected can be used as support regions for many 3D object recognition algorithms such as spin-images [2] or the tensor-based approach [4]. Figure 9 shows the curvedness of the ‘Buddha’ models at different scale levels r_k .

Algorithm 1 Multi-scale Curvedness-based Feature Extraction Algorithm

Data:

$\mathbf{P} = \{\mathbf{p}_i \in \mathbf{R}^3\}$: set of 3D points sampled from the surface.

$\mathbf{R} = \{r_k\}$: a set of scales.

Algorithm:

1: **for** $r \in \{r_k\}$ **do**

2: **for** $\mathbf{p} \in \{\mathbf{p}_i\}$ **do**

3: Find the neighbourhood N_r at scale r

4: Fit a jet to N_r

5: Compute principal curvatures k_1 and k_2

6: Compute the curvedness $c_{\mathbf{p}}$

$$c_{\mathbf{p}} = \sqrt{(k_1^2 + k_2^2)/2}$$

7: **end for**

8: Keypoints are positions \mathbf{p} having extremum values $c_{\mathbf{p}}$ both in the neighbourhood of radius r_k as well as over the above and below scales (r_{k-1}, r_{k+1}) .

9: **end for**

B. Feature Confidence

The *confidence value* of a feature located at a surface point \mathbf{p} at scale r_k is defined as

$$\gamma(\mathbf{p}, r_k) = \frac{|c_{\mathbf{p}} - \mu_{N_{\mathbf{p}}}|}{\sigma_{N_{\mathbf{p}}}} \quad (9)$$

where $N_{\mathbf{p}}$ is a set of all n 3D points in the neighbourhood of \mathbf{p} at not only scale r_k but also at the two adjacent scales r_{k-1} and r_{k+1} . $c_{\mathbf{p}}$ is the curvedness of \mathbf{p} as defined in (8). $\mu_{N_{\mathbf{p}}}$ and $\sigma_{N_{\mathbf{p}}}$ are the mean and standard deviation of the curvedness of all vertices in $N_{\mathbf{p}}$ respectively

$$\mu_{N_{\mathbf{p}}} = \frac{\sum_{\mathbf{p}_j \in N_{\mathbf{p}}} c_{\mathbf{p}_j}}{n} \quad (10)$$

$$\sigma_{N_{\mathbf{p}}} = \sqrt{\frac{\sum_{\mathbf{p}_j \in N_{\mathbf{p}}} (c_{\mathbf{p}_j} - \mu_{N_{\mathbf{p}}})^2}{n - 1}} \quad (11)$$

In this work, the reliability of a local feature is represented by its confidence value. If the confidence value is small, the feature may not be reliable because its curvedness value does not deviate far enough from the other values in the immediate neighbourhood. For example, if the confidence is small, the local maxima or minima may be selected as a result of measurement noise on the object's surface rather than a change in local surface structure. Assuming the distribution of the curvedness of a surface region is approximated by a normal distribution, about 68% of the curvedness values would be within $[\mu_{N_{\mathbf{p}}} - \sigma_{N_{\mathbf{p}}}, \mu_{N_{\mathbf{p}}} + \sigma_{N_{\mathbf{p}}}]$ and a threshold γ_t can be used to remove less reliable keypoints.

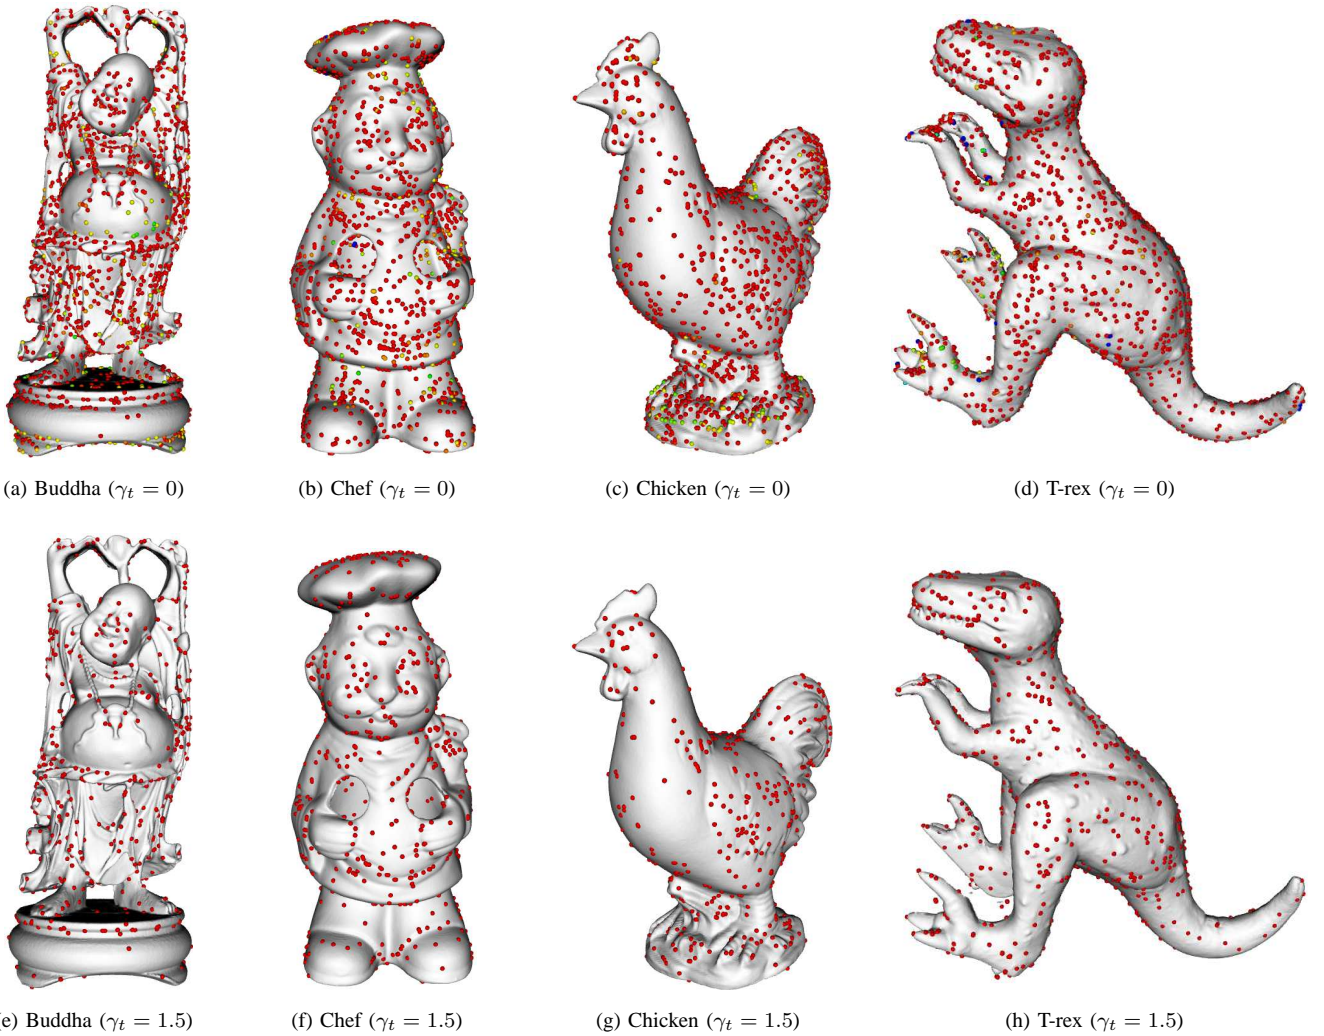


Fig. 10. Feature extraction results for four different 3D models with different confidence thresholds $\gamma_t = 0$ and $\gamma_t = 1.5$. Features are selected from each model using 7 different scales ($r = 1$ to $r = 7$).

V. EXPERIMENTS AND DISCUSSIONS

A. 3D Meshes

The proposed multi-scale feature extraction approach was tested on a variety of standard 3D models representing by triangular meshes. Figures 1 and 10 visualise the keypoints extracted from 5 different 3D models with two levels of threshold $\gamma_t = 0$ and $\gamma_t = 1.5$. The ‘Dragon’ and ‘Buddha’ together with many other models can be downloaded from the Stanford 3D Scanning Repository [8]. The ‘Chef’, ‘Chicken’ and ‘T-rex’ models were found at Mian’s website [22]. It can be seen from the figures that most of the salient positions in the models such as positions near the noses, mouths or eyes of the Buddha and Chef or the tail of the Chicken were selected as feature points.

Table I shows the comparison between the number of keypoints and the number of vertices in each model. It is clear from the table that the number of keypoints is significantly smaller than the number of vertices for all 5 surfaces. Without thresholding, the number of features is about 2% of the number of vertices in each model. When the threshold is set to $\gamma_t = 1.5$, the number of keypoints reduces to just about 0.5% of the number of vertices. Although the set of keypoints contains just a small percentage of the surface data, it is still a sparse but well-described representation of the geometric structures in the model as

TABLE I

COMPARISON BETWEEN THE NUMBER OF VERTICES AND NUMBER OF KEYPOINTS FOR ALL 5 MODELS WITH DIFFERENT CONFIDENCE THRESHOLDS.

Model	Number of vertices	Number of Keypoints ($\gamma_t = 0$)	Number of Keypoints ($\gamma_t = 1.5$)
Dragon	134559	2910 (2.16%)	927 (0.69%)
Buddha	133127	2760 (2.07%)	745 (0.56%)
Chef	176920	2439 (1.38%)	804 (0.45%)
Chicken	135142	2273 (1.68%)	740 (0.55%)
T-rex	176508	2910 (1.65%)	983 (0.56%)

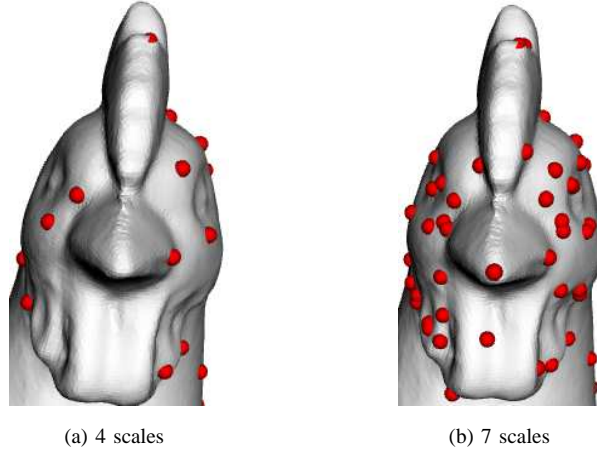


Fig. 11. Local features detected on the head of the ‘Chicken’ model using different numbers of scales in the scale-space representations.

evident in Figures 1 and 10.

B. Number of Scales

Figure 11 shows the keypoints detected on the head of the ‘Chicken’ model using two different numbers of scales in the scale-space representations. Two different numbers of scales, 4 and 7, are used in estimating the local features in Figure 11a and Figure 11b, respectively. It can also be seen from Figure 12 that the more scales used in the scale-space representation, the more features that are extracted and the better the geometric structures of the surface are represented by these features. One advantage of the multi-scale approach is that it can detect coarse-scale features even though the curvature might be low. We can see that important feature points such as the one on the tip of the nose can only be detected using a high number of scales. There is, of course, a trade-off between the number of scales and the time taken to extract local features from 3D surfaces. The more scales used, the more number of salient features can be detected. However, it is also more computationally expensive to process the scale-space representation of a 3D surface with too many scales.

C. Repeatability of Keypoints

It is very important that local features detected in the original surface will be present in the noisy data. In order to evaluate the repeatability of the feature points, white Gaussian Noise with standard deviation σ_G ranging from 0.001 to 0.1 was added to the 3D surfaces. When noise is introduced to the meshes, the variation of the local surface patches will increase. As a result, there would be more features points detected in noisy images compared to the original one. However, it is important that the local keypoints detected in the original surface will present in the noisy data.




Fig. 12. Number of features detected at different numbers of scales at two threshold levels $\gamma_t = 0$ and $\gamma_t = 1.5$ for the ‘Chicken’ model.




Fig. 13. Features detected from the face of the ‘Chef’ model with different noise levels.

Figure 13 shows the features extracted from the face of the ‘Chef’ model for different levels of noise. It can be seen from the figure that a large portion of local features from the original, smooth face are presented in the the noisy versions. For example, there are still many feature points lying around salient structures such as the nose, chin, eyes, mouth and ears even in the noisiest surface in Figure 13d. With the noise level of $\sigma_G = 0.005$, most of the keypoints in the original image appears in the noisy version. A quantitative evaluation of the repeatability of the features for five different 3D models is shown in Figure 14. At the noise level of $\sigma_G = 0.001$, nearly all of the features in the original model can be detected in the noisy surface. Even when the standard deviation of the noise goes to $\sigma_G = 0.1$, about 40% of the original features repeat in the noisy data.

D. Unstructured 3D Point Clouds

The proposed approach was also employed to extract local features from different unstructured point clouds without connectivity information between the vertices. When applying to 3D point clouds, Euclidean distance is used as the radius of the neighbourhood for collecting surrounding points. A kd-tree data structure is implemented to perform the local neighbour search efficiently. By fitting local manifolds directly to the surface points, the scale-space representation of the model can be created without the need to reconstruct the surface from the point cloud in advance which is a very error-prone and non-trivial process. Figure 15 shows the features extracted from two different point clouds of a truck and a tank. The sizes of the spheres represent the scales of the extracted features.




Fig. 14. Repeatability of the features for 5 different models in different noise conditions.

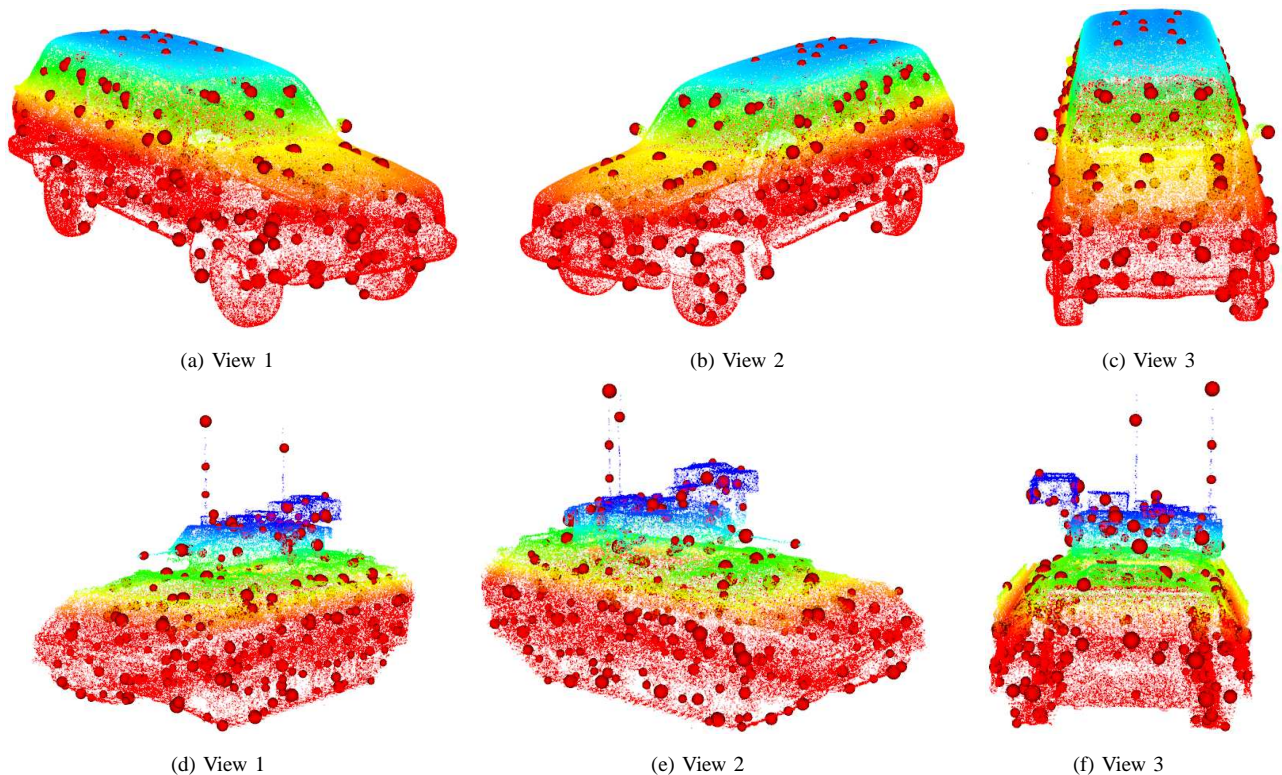


Fig. 15. Feature extraction results for two different unstructured point clouds, a truck and a tank, respectively. The size of each sphere denotes the scale at which the feature is selected. The point clouds are colour-coded using the z coordinates of the 3D points.

E. 3D Surface Registration using Spin-Images with Local Features

In this section, the results of combining our feature extraction framework and the spin-image algorithm [2] to register 3D surfaces are presented. The proposed method is used as a preprocessing step in order to improve the accuracy and efficiency of the spin-image registration algorithm. Due to the limitation of the spin-image implementation [2] in dealing with large datasets, each scan was re-sampled to contain about 10000 vertices using the cost driven approach proposed in [24]. To register two different views of the same object, we only performed the feature extraction on one view. The correspondences of spin-images

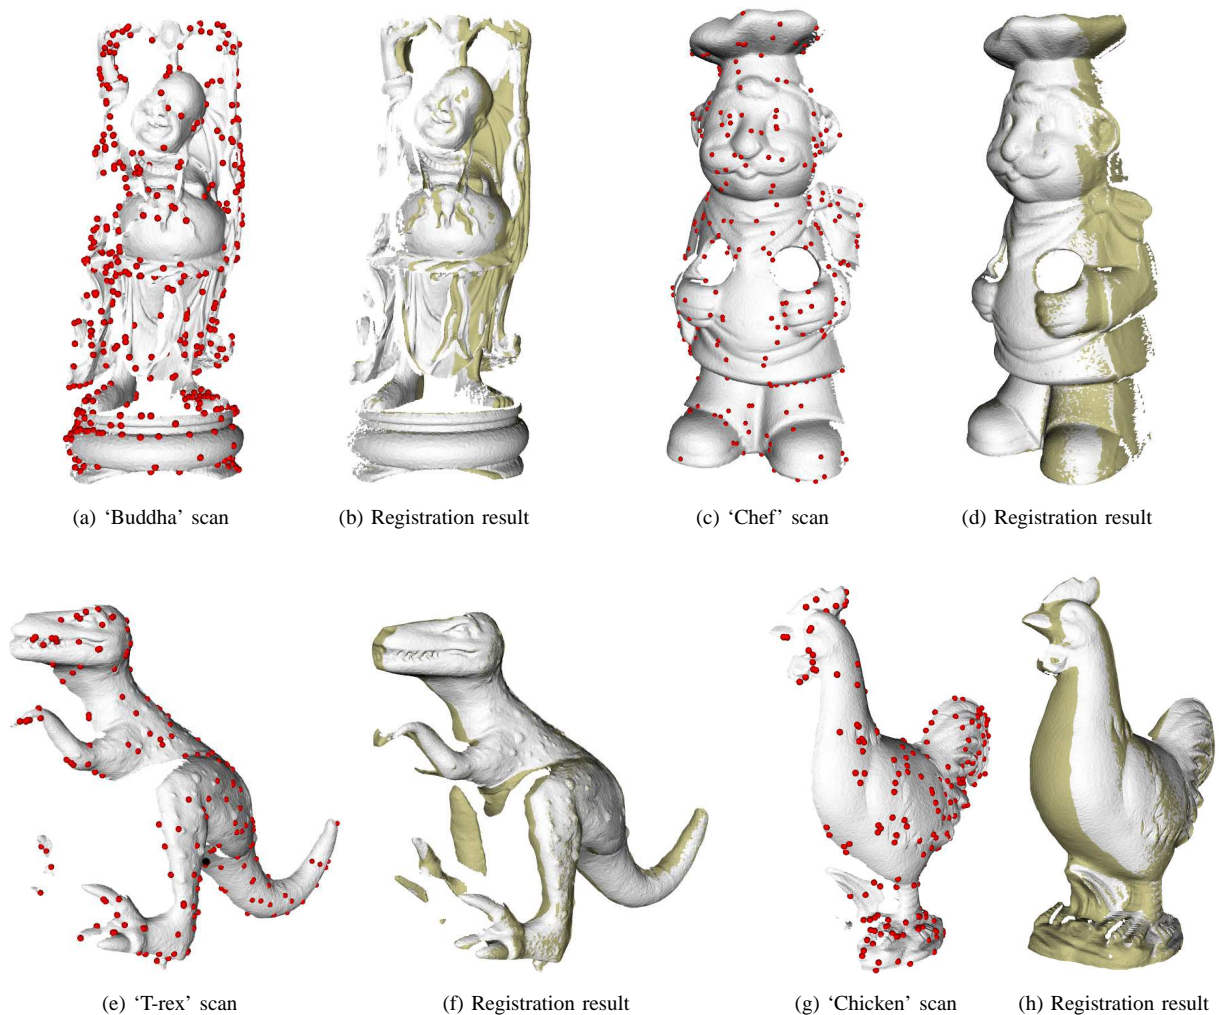


Fig. 16. Registration results for 2 different views of four 4D models. Features extracted from the model views are shown as red dots. Positions marked in darker color are the scene scan aligned w.r.t the model scan.

extracted at these features were search through all vertices of the other view to improve the probability of finding matches. Furthermore, the feature extraction step was performed off-line on the model scan so that it would not affect the actual matching time. Figures 16a, 16c, 16e and 16g show the features extracted from the model views of the ‘Buddha’, ‘Chef’, ‘T-rex’ and ‘Chicken’, respectively. As the difference between the registration results using random points and the ones using local features is not visually significant, only the registration results obtained by combining the spin-image algorithm and the proposed local feature extraction method are shown in Figures 16b, 16d, 16f and 16h.

A quantitative comparison of the registration results for two views of these models by using randomly selected points and our feature extraction approach is shown in Table II. It can be seen from the table that the number of features is much smaller compared to the number of randomly selected points. As a result, the registration time reduced significantly in all cases by employing the proposed feature extraction technique as a preprocessing step. The experiments were done on a Intel Core2Duo 2.4 GHz laptop with 2GB of memory running Linux. Besides the improvement in the speed of the matching process, our method also produced more accurate results as the average registration errors are smaller when using local features than if random surface points are selected. This error represents the average distance between all correspondences in a match after model points have been transformed by the match transformation [2]. The resolutions of the ‘Buddha’, ‘Chef’, ‘T-rex’ and

TABLE II
 QUANTITATIVE COMPARISON OF THE REGISTRATION USING RANDOMLY SELECTED POINTS AND LOCAL FEATURES.

Model	Random Points			Local Features		
	No. of selected points	Reg. time	Avg. error	No. of selected features	Reg. time	Avg. error
'Buddha'	6246	2m31s	0.965	370	5s	0.932
'Chef'	5371	2m54s	0.887	158	7s	0.812
'T-rex'	5183	2m34s	0.663	164	5s	0.607
'Chicken'	5163	2m56s	0.621	121	8s	0.616

'Chicken' meshes are 1.26, 2.06, 1.48 and 1.29, respectively.

VI. CONCLUSIONS

In this paper, a framework for extracting local features from 3D models using surface curvature has been presented. The scale-space representation of the surface geometry was constructed by estimating the local curvedness of the surface at different scales. By fitting a truncated Taylor expansion called the jet to the local surface patch, two principal curvatures as well as the curvedness can be approximated. Feature points were chosen as vertices on the surface having local extrema not only at the current scale but also at two adjacent scales. Furthermore, the reliability of each local feature was evaluated by using a confidence value that measures how far the curvedness of the feature is from the mean value of its neighbourhood. Experimental results on a number of different 3D models showed that the approach could robustly detect and localize local features from both 3D meshes and unstructured point clouds. Our method also appeared to work well in noisy conditions given the high repeatability of the features between the original and noisy range images.

In order to demonstrate the approach, the proposed feature extraction framework is used as a preprocessing step to the spin-image algorithm in order to perform the registration of 3D surfaces. It can be seen that both the accuracy and efficiency of the registration process are improved. In the future, we plan to extend our framework to handle other representations of 3D surfaces such as height maps, range maps or the recently proposed Canonical Face Depth Map (CFDM) [25]. Although the spin-images appeared to work well with our method, we will also be investigating the derivation of a new type of surface descriptor that uses not only the spatial distribution of surface points but also the curvature information to perform better recognition.

ACKNOWLEDGEMENTS

The authors would like to thank Prof. Doug Gray for his helpful comments on this work. We also would like to thank Ajmal S. Mian from the University of Western Australia and the Stanford Computer Graphics Laboratory for making their 3D databases publicly available. We are very grateful to Daniel Huber from the Robotics Institute, Carnegie Mellon University, USA for providing the spin-image implementation. The simulation data of 3D point clouds were provided by DSTO, Australia.

REFERENCES

- [1] A. Frome, D. Huber, R. Kolluri, T. Bulow, and J. Malik, "Recognizing Objects in Range Data using Regional Point Descriptors," in *Proc. European Conf. Computer Vision*, May 2004.

- [2] A. Johnson and M. Hebert, "Using Spin Images for Efficient Object Recognition in Cluttered 3D Scenes," *IEEE Trans. Pattern Analysis and Machine Intelligent*, vol. 21, no. 5, pp. 674–686, May 1999.
- [3] B. Matei, Y. Shan, H. Sawhney, Y. Tan, R. Kumar, D. Huber, and M. Hebert, "Rapid Object Indexing using Locality Sensitive Hashing and Joint 3D-Signature Space Estimation," *IEEE Trans. Pattern Analysis and Machine Intelligent*, vol. 28, no. 7, pp. 674–686, Jul. 2006.
- [4] A. Mian, M. Bennamoun, and R. Owens, "A Novel Representation and Feature Matching Algorithm for Automatic Pairwise Registration of Range Images," *IJCV*, vol. 66, no. 1, pp. 19–40, 2006.
- [5] R. Unnikrishnan and M. Hebert, "Multi-scale Interest Regions from Unorganized Point Clouds," in *Workshop on Search in 3D*, Jun. 2008.
- [6] X. Li and I. Guskov, "Multi-scale Features for Approximate Alignment of Point-based Surfaces," in *Proc. Eurographics Symp. on Geometry Proc.*, 2005.
- [7] J. Koenderink and A. van Doorn, "Surface Shape and Curvature Scales," *Image and Vision Computing*, vol. 10, pp. 557–565, 1992.
- [8] Stanford Computer Graphics Laboratory, "Stanford 3D Scanning Repository," <http://www-graphics.stanford.edu/data/3Dscanrep/>.
- [9] A. Witkin, "Scale-Space Filtering," in *Proc. IJCAI*, 1983, pp. 1019–1022.
- [10] T. Lindeberg, *Scale-Space Theory in Computer Vision*. Dordrecht: Kluwer, 1994.
- [11] D. Lowe, "Distinctive Image Features From Scale-Invariant Keypoints," *Int'l J. Computer Vision*, vol. 60, no. 2, pp. 91–110, 2004.
- [12] J. Koenderink, "The Structure of Images," *Biological Cybernetics*, vol. 50, pp. 363–370, 1984.
- [13] G. Taubin, "A Signal Processing Approach to Fair Surface Design," in *SIGGRAPH 95*, 1995.
- [14] M. Pauly, R. Keiser, and M. Gross, "Multi-scale Feature Extraction on Point-sampled Surface," in *Proc. Eurographics 2003*, 2003.
- [15] J. Novatnack and K. Nishino, "Scale-Dependent 3D Geometric Features," in *Proc. IEEE Int'l Conf. Computer Vision*, Oct. 2007.
- [16] A. Flint, A. Dick, and A. van den Hengel, "Thrifty: Local 3D Structure Recognition," in *Proc. DICTA 2007*, Dec. 2007, pp. 182–188.
- [17] M. Desbrun, M. Meyer, P. Schöder, and P. Barr, "Implicit Fairing of Arbitrary Meshes using Diffusion and Curvature Flow," in *SIGGRAPH 99*, 1999.
- [18] C. Dorai and A. Jain, "COSMOS: A Representation Scheme for 3D Free-Form Objects," *IEEE Trans. Pattern Analysis and Machine Intelligent*, vol. 19, no. 10, pp. 1115–1130, Oct. 1997.
- [19] F. Mokhtarian, N. Khalili, and P. Yuen, "Multi-scale Free-form 3D Object Recognition using 3D Models," *Image and Vision Computing*, vol. 19, pp. 271–281, 2001.
- [20] S. Yamany and A. Farag, "Surface Signatures: An Orientation Independent Free-form Surface Representation Scheme for the Purpose of Objects Registration and Matching," *IEEE Trans. Pattern Analysis and Machine Intelligent*, vol. 24, no. 8, pp. 1105–1120, 2002.
- [21] A. Jagannathan and E. Miller, "Three-dimensional Surface Mesh Segmentation using Curvedness-based Region Growing Approach," *IEEE Trans. Pattern Analysis and Machine Intelligent*, vol. 29, no. 12, pp. 2195–2204, 2007.
- [22] A. Mian, "A.S. Mian's 3D Models," <http://people.csse.uwa.edu.au/ajmal/3Dmodeling.html>.
- [23] A. Gray, *Modern Differential Geometry of Curves and Surfaces*. CRC Press, 1993.
- [24] P. Lindstrom and G. Turk, "Evaluation of Memoryless Simplification," *IEEE Trans. Visual. and Comp. Graph.*, vol. 5, pp. 98–115, 1992.
- [25] D. Colby and G. Stockman, "Real-time Identification using a Canonical Face Depth Map," *IET Computer Vision*, vol. 3, no. 2, pp. 74–92, 2009.