Pattern Recognition Letters 59 (2015) 41–47

Contents lists available at ScienceDirect

Pattern Recognition Letters

journal homepage: www.elsevier.com/locate/patrec

A simple and effective relevance-based point sampling for 3D shapes ☆

Emanuele Rodolà [a], Andrea Albarelli [b,*], Daniel Cremers [a], Andrea Torsello [b]

[a] Technische Universität München, Boltzmannstr. 3, Garching 85748, Germany
[b] Università Ca' Foscari Venezia, via Torino 155, Venice 30174, Italy

ARTICLE INFO

Article history:
Received 30 July 2014
Available online 6 April 2015

Keywords:
3D point sampling
Surface registration
Object recognition
Object in clutter

ABSTRACT

The surface of natural or human-made objects usually comprises a collection of distinct regions characterized by different features. While some of them can be flat or can exhibit a constant curvature, others may provide a more mixed landscape, abundant with high frequency information. Depending on the task to be performed, individual region properties can be helpful or harmful. For instance, surface registration can be eased by a set of non-coplanar smooth areas, while distinctive points with high curvature can be key for object recognition. For this reason, it is often critical to perform a surface sub-sampling that is suitable to the actual processing goal. To this end, most of the shape processing pipelines found in literature come bundled with one or more sampling rules, designed to boost their performance and accuracy. In this paper we introduce a sampling method for 3D surfaces that aims to be general enough to be useful for a wide range of tasks. The main idea of our method is to exploit the extent of the region around each point that exhibits limited local changes, granting higher relevance to points contained in compact neighborhoods. The effectiveness of the proposed method is experimented through its adoption as a point sampler within three very different shape processing scenarios.

1. Introduction

Point sampling is a key operation for many algorithms dealing with surfaces. Its adoption is needed for several reasons. If the surface to analyze is expressed as a parametric 3D curve, sampling is a useful discretization step to produce data which is easier to handle with standard algorithms. Even if the surface is represented as a triangulated mesh, sampling may help in reducing the total amount of points to handle, which can be mandatory if the complexity of the intended task is not linear and the meshes are large. However, the most important goal of sampling is probably the selection of surface points that are meaningful with respect to the task that is to be performed. To this end, it is quite natural that different sampling methods have been proposed to deal with each specific problem.

A quite standard example is the case of ICP surface registration [1,2]. This widely adopted method is used to obtain an accurate alignment between two coarsely registered surfaces. It is widely applied in the field of 3D scanning, where devices are able to capture only partial views and a proper alignment between them is needed to recover the complete surface of the scanned object. It basically works by iteratively minimizing a distance function measured between pairs of selected neighboring points. Regardless of the chosen distance function and matching criteria, for an accurate registration it is very important to sample points that are able to constrain well the rigid transformation between the processed surfaces. In fact, many different sampling variants have appeared in literature throughout the last decades. The normal space subsampling introduced by Rusinkiewicz and Levoy [3] attempts to sample uniformly on the sphere of normal directions rather than on the surface. The rationale is to avoid the predominance of large coplanar regions that would result in too many degrees of freedom. An interesting approach to better constrain the transformation is to select points that best equalize the error covariance matrix. To this effect Guehring [4] proposes to weigh the samples based on their contribution to the covariance matrix, but since the analysis is performed after the sampling, the approach cannot constrain the transformation if too few samples are chosen in a relevant region. On the other extreme, Gelfand et al. [5] propose an approach that selects the points that bind the transformation the most.

Sampling is also crucial for 3D object recognition tasks. However, differently from registration, the goal for the selected points set is not to be able to constrain a rigid motion, but rather to be distinctive enough to ease their recognition among different instances of the same object. The idea of point distinctiveness has been extensively used in image processing to develop interest point detectors such as the Harris Operator [6] and Difference of Gaussians [7]. While these approaches work well with 2D intensity images, they cannot be

☆ This paper has been recommended for acceptance by Dr. Y. Liu.

* Corresponding author. Tel.: +39 41 234 8465; fax: +39 41 234 8419.
E-mail address: albarelli@unive.it, andrea@albarelli.it (A. Albarelli).

http://dx.doi.org/10.1016/j.patrec.2015.03.009
0167-8655/

easily extended to handle 3D surfaces since no intensity information is directly available. Several efforts have been made to use other local measures, such as curvature or normals to find relevant points on a surface. One of the first descriptors to capture the structural neighborhood of a surface point was described by Chua and Jarvis, who with their Point Signatures [8] suggest both a rotation and translation invariant descriptor and a surface matching technique.

More recent 3D interest point detectors include Harris 3D [9], a generalization of the Harris 2D detector to Euclidean surfaces, Normal Aligned Radial Features [10], making explicit use of object boundary information, and Intrinsic Shape Signatures [11], providing a weighted occupational histogram of data points, computed with respect to a local intrinsic 3D reference frame. In order to guarantee such frame to be stable, it is aligned on selected salient features characterized by large three dimensional point variations. Such property is assessed by looking at the smallest eigenvalue of the point scatter matrix of the feature neighborhood. An additional check is finally performed on the ratios of the eigenvalues to avoid ambiguous frames resulting from symmetries. Zaharescu et al. [12] presented an approach for feature point detection (MeshDOG) and description (MeshHOG), based on the value of any scalar function defined over the surface (i.e., curvature or texture, if available). Other widely adopted 3D point descriptors include Spin Images [13], a rich characterization obtained by a binning of the radial and planar distances of the surface samples respectively from the feature point and from the plane fitting its neighborhood, and SHOT [14], which, in addition to the feature vector, also derives a repeatable local reference frame. The interested reader can refer to [15] for a comprehensive evaluation of recent 3D keypoint detectors.

In this paper we introduce a general sampling method, described in the next section, which aims to select points that can be successfully adopted within all the described scenarios. To this end, we associate to each surface point a weight, named *relevance*, assessing its degree of uniqueness with respect to the region in which it is contained. The basic idea is that the relevance should be high for points that have unique normal orientation with respect to their surroundings, while it should be low for evenly oriented patches. Also, relevance should be inversely proportional to the area of the neighborhood, thus fostering points that do not belong to large regions of uniform curvature. Furthermore, the measure should be computed through an integral measure, thus making it robust with respect to noise. Such relevance can be used to define the probability density distribution upon which the actual sampling is based. The rationale of a relevance-based sampling is that distinctive points that can be adopted for object recognition usually correspond to ridges, corners or valleys. Such features will obtain high relevance values, and thus should be favored in the selection. At the same time, points lying in flat areas, while yielding low relevance, can still be selected due to their large number. Also, areas with uniform curvature are expected to be regularly sampled over all their span in a similar manner to what would happen adopting normal space subsampling. The stated all-roundness of this sampling has been put to the test in the experimental section, where we use it as a drop-in replacement for several state-of-the-art points selection methods within three tasks.

2. Contribution and application scenarios

The goal of our method is to introduce a general purpose sampler that yields points that can be adopted successfully to solve problems ranging from surface registration to object recognition. To this end, our sampler is not to be considered an interest point detector and the relevance measure we are introducing cannot be directly translated into a distinctiveness assessment. In fact, while distinctive points can be paramount for object recognition or classification, they could not be enough distributed over the surface for accurate fine alignment. Differently, our relevance measure and sampling schema aim at the selection of characterizing points scattered over all the targeted shape, accounting for both distinctiveness and coverage. The first goal is obtained by giving to distinctive points higher sampling probability. The second objective is reached by making relevance inversely proportional to the area of flat regions. This way, while single points in large homogeneous areas exhibit low sampling probability, their large amount still make it possible for a number of them to be selected.

While many task-specific detectors and samplers can be found in literature, we feel that the use cases for a general purpose surface samplers are several, especially within the increasingly important context of large and distributed databases of 3D objects and surfaces. For instance, it would make sense to be able to extract from each shape stored a reduced number of representative points that can be adopted for different tasks without ever needing to access the original data. Of course, it would be unreasonable to expect such a general purpose set of characterizing points to outperform every specially crafted selector. Still, we will show that it can be used as a sound alternative in many scenarios, scoring comparable or better results than task-specific approaches.

3. Relevance-based sampling

Central to our sampling strategy is the concept of *relevance* defined for each point over a mesh. The *relevance* of a point p is related to how similar points around p are to it. The larger the number of similar points, the less distinctive, and thus the less relevant, p is. For this reason we formalize the idea of relevance of point p in terms of the area of a surface patch around it where points have a similar orientation. This is a very simple similarity notion, as it only accounts for the point normals which, when not already available, can be easily estimated on any mesh. Despite its crudeness, it captures important aspects of the surface, such as local curvature and structure. Moreover, it facilitates the distribution of the relevance (and thus of the sampling) over all the surface orientations, akin to normal space samplers. According to these considerations, the computation of the relevance for a point is strictly connected to the formalization of *influence area*:

Definition 1 (Influence Area)**.** Let p be a point of surface S, we associate to it an *Influence Area* A_p such that

$$A_p = \{q \in S | N_p^T N_q > T \text{ and } p \sim q\} \tag{1}$$

where N_p and N_q are the normals of the surface S at points p and q, while $p \sim q$ means that there is a path in A_p connecting p to q, and the dot threshold T is a parameter of the approach.

For small values of T the area of A_p is related to the average absolute radius of curvature

$$||A_p|| \approx \bar{r} = \frac{|r_1| + |r_2|}{2} = \frac{|1/k_1| + |1/k_2|}{2}, \tag{2}$$

where $||A_p||$ denotes the area of region A_p, k_1 and k_2 are the principal curvatures of S in p and $r_1 = 1/k_1$ and $r_2 = 1/k_2$ are the radii linked with the principal curvatures. Points within A_p are well aligned to the normal of p and if the surface orientation varies quickly in one direction the growth of the region in that direction will be limited, thus the size of A_p is linked with the distinctiveness of p. The area will be inversely proportional to curvature, along edges it will extend only in one dimension attaining a size one order of magnitude smaller, and will be almost point-like on vertices, where surface alignment would be locally completely constrained with the exception of rotations around the point normal. Hence, the area is inversely proportional to how much the surface is constraining the transformation locally, making the method suitable also for registration.

A representation of the expected growth of the influence area for the three aforementioned scenarios is shown in Fig. 1. With the first example the point p is placed on the top of an isotropically smooth region and the area A_p extends symmetrically in all the directions from