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, 85748 Garching, Germany*
[b]*Università Ca' Foscari Venezia, via Torino 155, 30174 Venice, Italy*

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 (Chen and Medioni, 1991; Besl and McKay, 1992). 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 (2001) 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 (2001) 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. (2003) propose an approach that selects the points that bind the transformation the most.

[**]Corresponding author: Tel.: +39-41-234-8465; fax: +39-41-234-8419
  e-mail: `albarelli@unive.it` (Andrea Albarelli)

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 (Harris and Stephens, 1988) and Difference of Gaussians (Marr and Hildreth, 1980). While these approaches work well with 2D intensity images, they cannot be 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 (Chua and Jarvis, 1997) suggest both a rotation and translation invariant descriptor and a surface matching technique.

More recent 3D interest point detectors include Harris 3D (Sipiran and Bustos, 2010), a generalization of the Harris 2D detector to Euclidean surfaces, Normal Aligned Radial Features (Steder et al., 2011), making explicit use of object boundary information, and Intrinsic Shape Signatures (Zhong, 2009), 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. (2009) 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 (Johnson and Hebert, 1999), 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 (Salti et al., 2014), which, in addition to the feature vector, also derives a repeatable local reference frame. The interested reader can refer to Tombari et al. (2013) 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

Fig. 1. Region A_p grows in all directions in a smooth part of the surface, in only one direction along edges and boundaries and grows little on vertices.

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*

$$A_p = \{q \in S \mid N_p^T N_q > T \text{ and } p \sim q\} \qquad (1)$$

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

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

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 Figure 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 the point p until the condition expressed by Equation 1 is not satisfied anymore. In the second example p is placed over a ridge and A_p is allowed to grow only along the axis characterized by a small curvature, while it stops quite abruptly along the orthogonal direction. Finally, in the third example, p has been chosen on the top of a spike, hence the growth of A_p is severely limited along all the directions radiating from p.

With the patches A_p at hand, we can compute for each point p the measure of relevance that will be used within our sampling schema.

Definition 2 (Relevance). *Given the influence area A_p of the point p, its relevance can be defined as:*

$$f(p) = \|A_p\|^{-k} \qquad (3)$$

where k is an equalization parameter, changing the relative weight of "common" and "distinctive" point. In particular, the larger the value of k, the more the distinctiveness of points forming a small patch A_p is emphasized.

Since the region A_p is defined in terms of an angular threshold, $\|A_p\|$ is invariant with respect to resampling, up to the precision imposed by the new sampling resolution. Further, any scale change varies the areas proportionally, so the ratio between patch areas is scale-invariant. Finally, the area of A_p is an integral measure, thus being less sensitive to noise, and varies continuously along the surface, with T being a smoothing factor. When the surface is discretized into points and edges, A_p can be easily computed with a region growing approach starting from each point p. If the regions are big, one could use the continuity and locality of A_p to update the region from neighboring points, but in practice, we add a size threshold D limiting the growth of A_p to points whose distance from p is less than D. This way we limit the complexity of the region growing process to $O(D^2)$ for each point and we avoid the uncontrolled expansion of A_p on flat surfaces.

Note also that, out of practical reasons, we can *approximate* the measure of relevance $f(p)$ for each point p with $\hat{f}(p)$:

$$\hat{f}(p) = |A_p|^{-k} \qquad (4)$$

where $|A_p|$ is the number of points in A_p. This approximation works under the assumption that the edge length is uniform through the discretization of surface S. If this is not the case, a more accurate approximation can be obtained as the sum of the Voronoi area of all the points included in A_p.

The effect of the described parameters is qualitatively shown using a JET color scale in Figure 2. Each test has been performed by changing one parameter and keeping fixed the others. In the first row, we test the effect of the dot product threshold T in equation (1). The lower threshold (corresponding to about 30 degrees) leads to a too wide region expansion, thus assigning low relevance to most points. By contrast, the larger value (about 8 degrees) results in an exceedingly sharp distribution and is also too sensitive to noise, as small perturbations may break the region growing early. A medium value for T seems to yield the best results as it allows to assign high relevance to relevant points that lie in distinctive regions. Notice also that the relevance map obtained with $T = 0.93$ mostly favors distinctive points located within high curvature areas, while a slightly larger value ($T = 0.96$) captures more points

Fig. 2. Effect of the parameters over the relevance of surface points on the Bunny model (see text for detail).

distributed all over the surface. The first distribution could be more suitable for recognition tasks, where strong distinctiveness is important, while the more uniform coverage of the second might help with registration. To this end, parameter T can be considered a useful tuning knob that can be turned according to the problem the sampling will be used for. In the second row, we perform the same analysis with respect to parameter k. Differently from T, k has no effect in the area computation, but just in how areas are converted to relevance. In practice, this parameter acts as a regulator of relevance steepness, making it more uniform for small values and increasing its entropy for high values. When k is low, most points exhibit similar relevance, when it is high only a few very distinctive points maintain non-zero relevance. In most applications a fit value for k is the one that guarantees a relevance map which is both equalized and well distributed within its range (to this end, values around 1 seem to behave reasonably). Finally, we show the effect of parameter D in the third row (expressed in edge resolution units). Since the only role of D is to stop the computation of A_p earlier, it has a minimal impact on the relevance computation, as long as it is not too low, in which case the overall process can be too much subject to noise (for instance with $D = 2$ in Figure 2).

Once we have computed the distinctiveness of all points on the surface, we can proceed to sample points with a density proportional to $\hat{f}$. To do this we select any ordering $p_1, \ldots, p_n$ of the points in S and compute the cumulative function

$$\hat{F}(p_i) = \sum_{j=1}^{i} \hat{f}(j), \tag{5}$$

then we sample a number x uniformly in $[0, \hat{F}(p_n)]$ and find the smallest index i such that $\hat{F}(p_i) > x$. To perform the search effi-

ciently, we use interpolation search (Perl et al., 1978), a variant of binary search that instead of splitting the interval $[i, j]$ in half at each iteration, it splits it at point $i + \frac{x - \hat{F}(p_i)}{\hat{F}(p_j) - \hat{F}(p_i)}(j - i)$. It is a well known result that interpolation search finds an element in a sorted array in $O(\log \log n)$ on average for near-uniformly distributed data, compared to the $O(\log n)$ complexity of binary search. Further, the search is faster the higher the entropy of f is. This results in an expected $O(m \log \log n)$ complexity when sampling m relevant points from a surface containing n points.

In Figure 3 we present the cumulative distribution function (CDF) over the Bunny model corresponding to the four values of T shown in the first row of Figure 2. Each CDF has been obtained by normalizing $\hat{F}(p_i)$ by dividing it for its maximum value. Higher values for T result in steeper curves, making the whole process more selective.
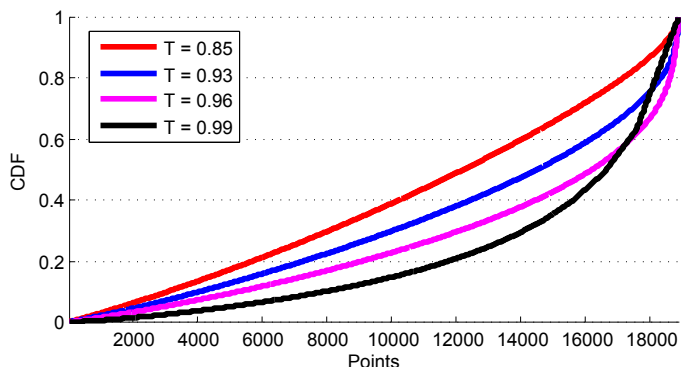

Fig. 3. CDF for several values of the threshold T.

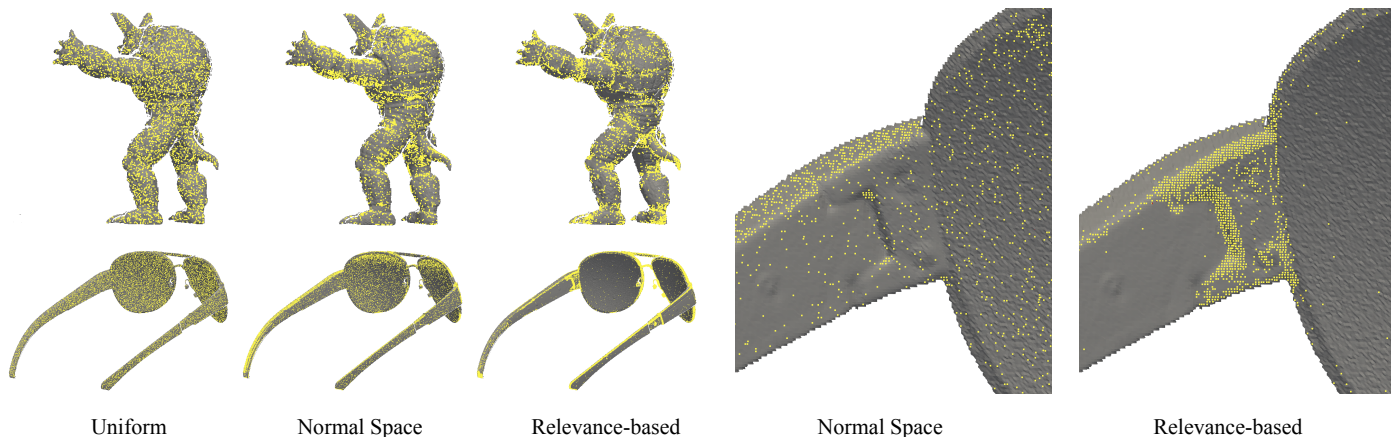| Uniform | Normal Space | Relevance-based | Normal Space | Relevance-based |

Fig. 4. Examples of the different sampling approaches and close-ups. relevance-based sampling concentrates samples along the surfaces' fine structures.

4. Experimental Evaluation

According to our stated goals, we designed a set of experiments specially crafted to assess the performance of the proposed approach with respect to three different problems. Namely, we tested the relevance-based sampling when used as a supplier of samples for the following tasks:

- Surface registration with ICP (Rusinkiewicz and Levoy, 2001). Relevance-based sampling is adopted as a substitute for uniform sampling and normal space subsampling;
- Object recognition in cluttered scenes (Rodolà et al., 2013). Our method is used as a replacement for the sampler proposed in the paper;
- Shape classification based on a Bags of Visual Words approach (Csurka et al., 2004). Relevant samples are extracted in order to compute distinctive descriptors.

All the experiments have been performed with a value for the parameter D of 6 times the average edge resolution.

4.1. Fine Surface Registration through ICP

To evaluate the performance of the sampling approach we created several range images with known ground-truth transformations. To this end we took the 3D models of the Bunny, Armadillo, Dragon, and Buddha models from the Stanford 3D scanning repository, and range maps extracted from six sets of glasses scanned using a home-brew scanner built in our lab. Figure 4 shows an example of sampling using the three tested strategies. In the experiments we used 18 scans for each model. For the glasses we used directly the range images provided by the scanner, while for the models taken from the Stanford repository the range images were created by projecting the models onto virtual orthographic cameras placed on a ring around them. Once the range images were to hand, additive Gaussian noise was added along the z dimension to simulate measurement error. In order to avoid having perfect point correspondences, the virtual shots, and thus the points in the various range images, were obtained by projecting equally spaced points on the view-plane of the virtual cameras, and the depths were com-
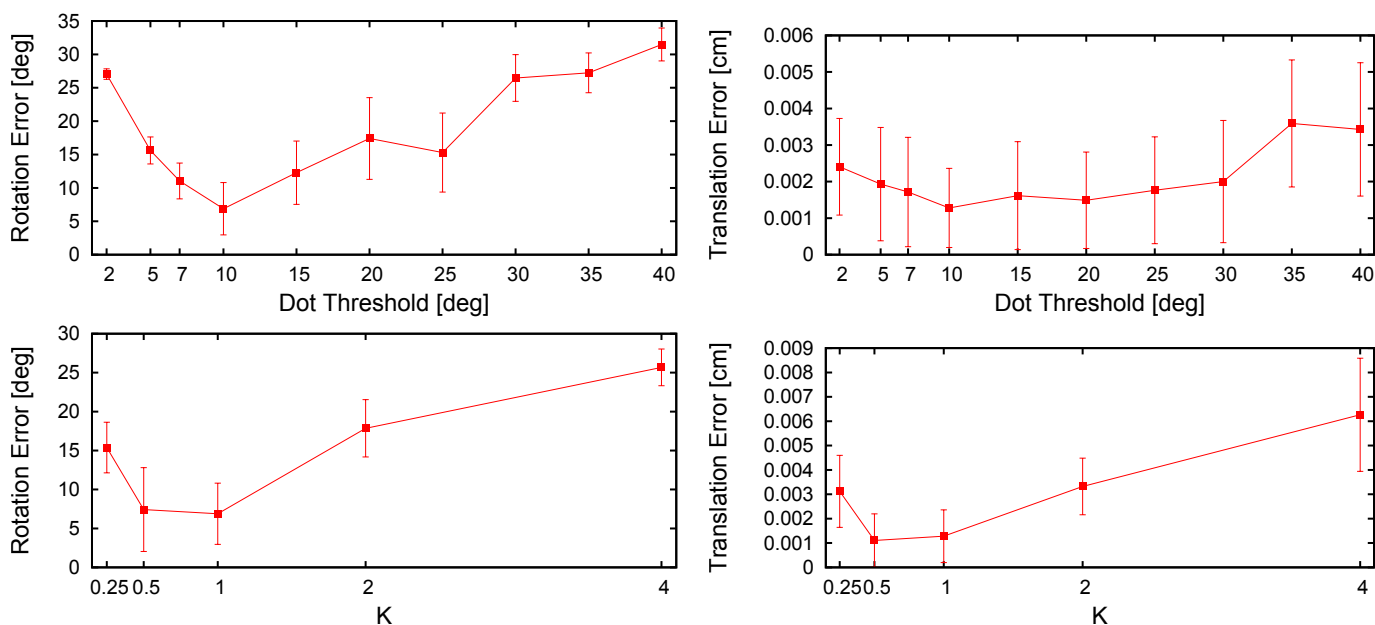

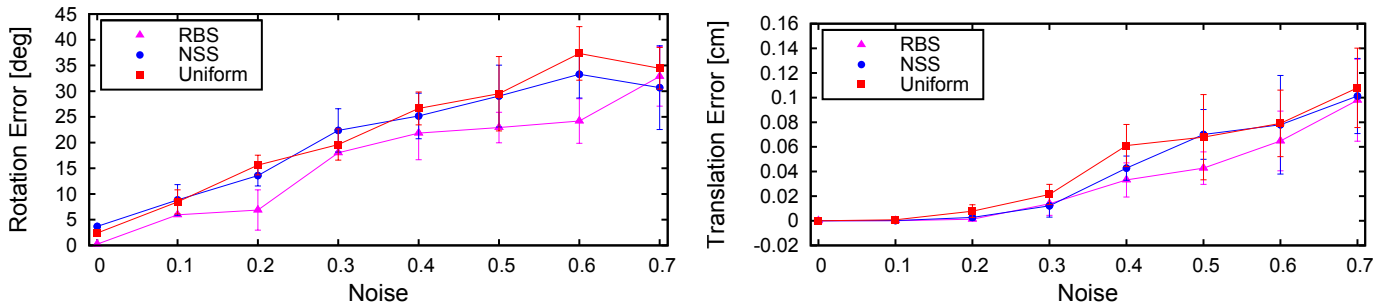**Fig. 5.** Effects of parameters on rotational and translational error.

Fig. 6. Comparison of rotational and translational error obtained with the three sampling strategies.

puted by finding the first intersection of the rays with the model. All the measures of quality of the alignments are based on the ground-truth alignment, and not the usual Root Mean Square Error (RMSE) because the value of the RMSE depends heavily on the sampling strategy and it is completely blind with respect to the noise overfitting problem. All the surface registration tests have been performed using a sampling of 1000 points from each surface for each one of the tested samplers.

The first set of quantitative experiments is a sensitivity analysis trying to assess the role of the dot threshold T and of the equalization parameter k. Figure 5 shows the angular error (in degrees) and the translation error (in centimeters) as a function of the two parameters. In all the cases the pairs of range images were selected randomly at a distance along the view-circle of at most 3 positions (90 degrees distance in view direction) and were perturbed with additive Gaussian noise along the z dimension with standard deviation equal to 0.4 times the average edge length. We can clearly see that there is an optimal value for the dot threshold at around 10 degrees, and it appears that the optimal value for the equalization parameter k is just slightly below 1. This is due to the fact that noise affects the size of small regions more than larger ones, keeping them smaller, resulting in over-inflated relevance values. A value of k smaller than 1 balances this effect by reducing the relative weight of the smaller regions with respect to larger ones.

Finally, Figure 6 plots the resulting rotation and translation error of the alignments obtained with the three sampling strategies as a function of the level of noise added to the range images. We can see that the relevance-based sampling consistently outperforms uniform sampling by a large margin in both rotational and translational error. Normal space sampling, on the other hand, has the same performance as uniform sampling for rotational errors, while it exhibits the same low translational error obtained by the relevance-based sampling for noise levels smaller than 0.4 times the average edge length. This is consistent with the fact that normal space sampling constrains only the translational sliding. Note however, that with larger noise levels the translational error of normal space sampling breaks down to uniform sampling levels. This is probably due to the interaction between bin-size and noise, with high noise spreading neighboring points onto several bins.

4.2. Object Detection in Cluttered Scenes

One of the most recent techniques for model-based object detection in general 3D scenes has been introduced in Rodolà

et al. (2013) and Albarelli et al. (2011). This approach allows to search for a model shape within a cluttered scene that can contain a noisy and possibly incomplete version of the model. This is the case, for instance, when searching for a specific object in range images or other 2.5D imagery. The method does not use descriptors for matching, as it is based on spatial relations between points. However, in order to reduce the complexity and to enhance the chance of performing correct recognition, a proper surface sampling is performed as a preliminary step. Specifically, an Integral Hash is computed over all the points of the mesh (Albarelli et al., 2010). Afterwards, only points with a value of the integral hash above a given threshold are retained as significant and thus used for matching. For this set of experiments we ditched the integral hash sampling and we adopted the proposed relevance-based method, leaving the remaining parts of the pipeline unaltered. Since the original game-theoretic method already obtained very good results for the tested dataset, we restricted our experiments to the 30 most challenging scenes (i.e. the ones that yielded the largest number of failures in the original paper, see Figure 7 for an example).

Specifically, we tested our relevance measure by using three different sampling schemas:

- **Relevance Based Random Sampling (RBRS):** we adopted the random sampling described in Section 3. We sampled a total of 1000 points;

- **Relevance Based High Thresholding (RBHT):** we computed the relevance for each point and we set a threshold for point acceptance so as to obtain exactly the 1000 points with higher relevance;

- **Relevance Based Low Thresholding (RBLT):** as for the pevious case we are selecting points using a hard threshold, this time setting it to be low enough to retain exactly 2 times 1000 points.


**Fig. 7. A recognized model in a very challenging scene.**

Fig. 8. Comparison between the results obtained with our method used as a sampler and as a thresholded point selector for different values of parameters T and k. Other baseline methods have been added as a reference (see text for details). The same legend applies to both plots.

The rationale for testing these three sampling schemes is to study if the relevance can be used directly as a point selector without losing any performance or if, differently, a relevance-driven random sampling is a preferable choice also for object recognition tasks. Notice that a relevance-based hard thresholding is just a sampling rule that assigns a probability to be extracted of exactly 1 to points above the threshold and of exactly 0 otherwise. In addition to these three relevance-based point samplers we included in the comparison three other baseline distinctive point selectors: the original integral hashes (IH) adopted in (Rodolà et al., 2013), the Intrinsic Shape Signatures (ISS) and the Harris 3D (H3D). Again, for each case we extracted a total of 1000 points.

In Figure 8, we show the recognition rates achieved using RBRS, RBHT and RBLT by respectively changing values of T (with a value of $k = 1$) and of k (with a T corresponding to 15 degrees). All the results have been obtained using the adaptive variant of (Rodolà et al., 2013). The other methods appear in both plots as horizontal lines since they do not depend on parameters k and T. For IH we used the best parameterization presented in the original paper and we tested separately the results obtained using the adaptive and fixed technique.

The best performance happens with RBRS when $k = 1$ and T sets a threshold angle between 15 and 20 degrees. This is consistent with the observation made in Section 3. In fact, as seen in Figure 2, a slightly lower value for T (corresponding to a wider threshold angle), concentrates the relevance (and thus the sampling) on more distinctive areas, supporting better recognition rates. Still, even if relevance seems to be important, RBHT is unable to obtain comparable results with the same number of samples used by RBRS. This is probably due to the fact that the recognition method used has better chances of finding rigid-enforcing correspondences when the points are distributed over the whole surface, rather than segregated only within few very relevant areas. This is somehow confirmed by the fact that RBLT performs better and obtains a performance level that is comparable with RBRS. However, this comes at the price of a doubled number of samples which, in turn, results in a much larger convergence time for the recognition, since the adopted method is quadratic with the number of points. Finally, we can observe that the proposed technique outperforms also the baseline methods IH, ISS and H3D.

4.3. Descriptor-based 3D Shape Classification

The third task for which we tested the suitability of the proposed sampling method is the classification of 3D shapes. To this end, we adopted a baseline classification method based on the so-called *bag of visual words* approach as implemented in the VLFeat library (Vedaldi and Fulkerson, 2010). Specifically, we used the Wave Kernel Signatures (WKS) (Aubry et al., 2011) to produce data vectors from surface points. The set of points to be used to compute WKS descriptors (and thus for classification), are selected using our method, uniform random subsampling and three different 3D keypoint detectors, respectively the Harris 3D (Sipiran and Bustos, 2010), the Intrinsic Shape Signatures (Zhong, 2009) and the Heat Kernel Signature (Sun et al., 2009). For the former two, we relied on the implementations provided by the PCL library (Rusu and Cousins, 2011).

We adopted the TOSCA dataset (Bronstein et al., 2008), considering 7 classes of objects undergoing non-rigid deformations (cats, dogs, horses, centaurs, gorillas, female figures, male figures). With respect to the parameters for our method, we used $k = 1$ and $T = 0.96$ (about 16 degrees). This is due to the fact that such values have proven to be a reasonable choice for recognition tasks, which suggests that they should be suitable also for classification, where distinctiveness and repeatability are certainly sought. In Figure 9 we show, for each class, the percentage of correct classifications obtained using the five sampling methods. These values are obtained by training the classifier using two randomly selected shapes and by testing it on four validation objects. The experiment is repeated several times with different training and validation sets. From the graph it is apparent that our method allows to obtain consistently better results with respect to uniform sampling, Harris 3D and Heat Kernel Signatures. Still, Intrinsic Shape Signatures exhibit a higher ratio of correct classification for a couple of classes. To obtain a global comparison considering all the classes, we created an additional meta-class labelled *Average* where the true-positive rate is computed as the average over all the 7 original classes. This latter test shows a slight advantage in using the proposed relevance-based sampling method.

4.4. Execution time

As a final evaluation we compared the execution times for the sampling step using either the binary search or the interpo-
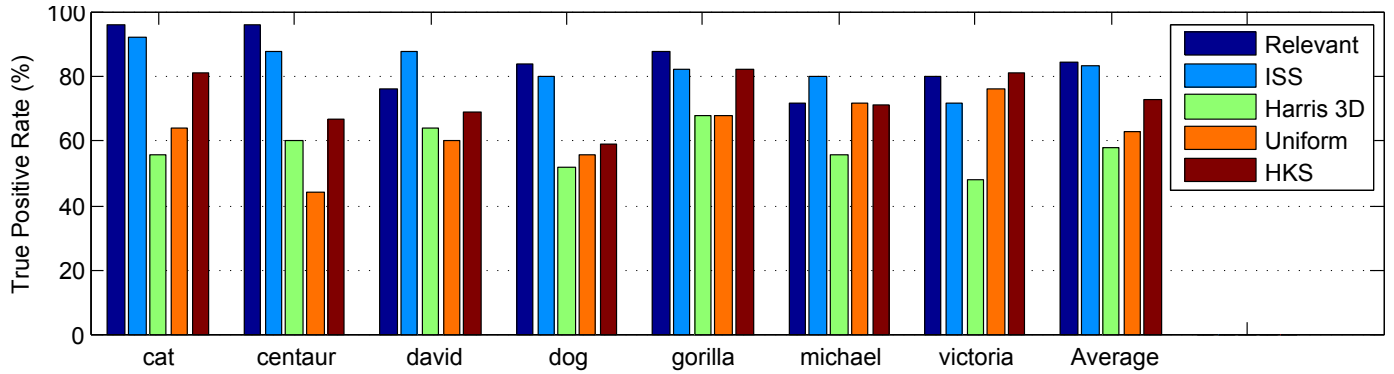
Fig. 9. Percentage of correct classifications for each class using the five tested sampling methods.

lation search as suggested in Section 3. We considered only the sampling step since the total running time depends also on the task-specific part of the algorithm. Average results on models of ~20k vertices are shown in the following table:

Method	500 pts	5k pts	10k pts
Binary search	0.06s	1.08s	2.85s
Interpolation search	0.06s	0.89s	1.14s

While the advantage of interpolation search is not huge, we still think that the required implementation effort is low enough to fully justify its adoption.

5. Conclusions

In this paper we have presented a novel surface sampling strategy, based on the local distinctiveness of each point. Such distinctiveness is gauged through an integral measure, called relevance, that is robust with respect to noise. The points are then sampled with a density proportional to their distinctiveness. The method is easy to implement and it is genuinely adaptable, as the selectivity of the sampling can be easily tuned through parameters. Since we are proposing our sampling as a general-purpose tool, we explicitly tested its suitability to deal with some classical tasks: surface registration, object detection in cluttered scenes, and 3D shape classification. To this end, we consistently obtained competitive results, comparable or even outperforming sampling methods that were specifically designed for each task.

References

Albarelli, A., Rodolà, E., Bergamasco, F., Torsello, A., 2011. A non-cooperative game for 3d object recognition in cluttered scenes, in: Proceedings - 2011 International Conference on 3D Imaging, Modeling, Processing, Visualization and Transmission, 3DIMPVT 2011, pp. 252–259.

Albarelli, A., Rodolà, E., Torsello, A., 2010. Loosely distinctive features for robust surface alignment, in: Daniilidis, K., Maragos, P., Paragios, N. (Eds.), Computer Vision – ECCV 2010. Springer Berlin Heidelberg. volume 6315 of *Lecture Notes in Computer Science*, pp. 519–532.

Aubry, M., Schlickewei, U., Cremers, D., 2011. The wave kernel signature: A quantum mechanical approach to shape analysis, in: Proc. ICCV Workshops, pp. 1626–1633.

Besl, P.J., McKay, N.D., 1992. A method for registration of 3-d shapes. IEEE Trans. Pattern Anal. Mach. Intell. 14, 239–256.

Bronstein, A., Bronstein, M., Kimmel, R., 2008. Numerical Geometry of Non-Rigid Shapes. 1 ed., Springer Publishing Company, Incorporated.

Chen, Y., Medioni, G., 1991. Object modeling by registration of multiple range images, in: Proc. ICRA, pp. 2724–2729 vol.3.

Chua, C.S., Jarvis, R., 1997. Point signatures: A new representation for 3d object recognition. Intl. J. of Comput. Vis. 25, 63–85.

Csurka, G., Dance, C.R., Fan, L., Willamowski, J., Bray, C., 2004. Visual categorization with bags of keypoints, in: In Workshop on Statistical Learning in Computer Vision, ECCV, pp. 1–22.

Gelfand, N., Ikemoto, L., Rusinkiewicz, S., Levoy, M., 2003. Geometrically stable sampling for the ICP algorithm, in: Int. Conf. 3-D Digital Imaging and Modeling.

Guehring, J., 2001. Reliable 3d surface acquisition, registration and validation using statistical error models, in: Int. Conf. 3-D Digital Imaging and Modeling.

Harris, C., Stephens, M., 1988. A combined corner and edge detector, in: Proc. Fourth Alvey Vision Conference, pp. 147–151.

Johnson, A.E., Hebert, M., 1999. Using spin images for efficient object recognition in cluttered 3d scenes. IEEE Trans. Pattern Anal. Mach. Intell. 21, 433–449.

Marr, D., Hildreth, E., 1980. Theory of edge detection. Royal Soc. of London Proc. Series B 207, 187–217.

Perl, Y., Itai, A., Avni, H., 1978. Interpolation search – a log log N search. CACM 21, 550–553.

Rodolà, E., Albarelli, A., Bergamasco, F., Torsello, A., 2013. A scale independent selection process for 3d object recognition in cluttered scenes. International journal of computer vision 102, 129–145.

Rusinkiewicz, S., Levoy, M., 2001. Efficient variants of the icp algorithm, in: Third Intl. Conf. on 3D Digital Imaging and Modeling, pp. 145–152.

Rusu, R., Cousins, S., 2011. 3d is here: Point cloud library (pcl), in: Robotics and Automation (ICRA), 2011 IEEE International Conference on, pp. 1–4.

Salti, S., Tombari, F., Stefano, L.D., 2014. Shot: Unique signatures of histograms for surface and texture description. Computer Vision and Image Understanding 125, 251 – 264.

Sipiran, I., Bustos, B., 2010. A robust 3d interest points detector based on harris operator, in: Proceedings of the 3rd Eurographics Conference on 3D Object Retrieval, Eurographics Association, Aire-la-Ville, Switzerland, Switzerland. pp. 7–14.

Steder, B., Rusu, R., Konolige, K., Burgard, W., 2011. Point feature extraction on 3d range scans taking into account object boundaries, in: Robotics and Automation (ICRA), 2011 IEEE International Conference on, pp. 2601–2608.

Sun, J., Ovsjanikov, M., Guibas, L., 2009. A concise and provably informative multi-scale signature based on heat diffusion, in: Proc. SGP, Eurographics Association. pp. 1383–1392.

Tombari, F., Salti, S., Di Stefano, L., 2013. Performance evaluation of 3d keypoint detectors. International Journal of Computer Vision 102, 198–220.

Vedaldi, A., Fulkerson, B., 2010. Vlfeat: An open and portable library of computer vision algorithms, in: Proceedings of the International Conference on Multimedia, ACM, New York, NY, USA. pp. 1469–1472.

Zaharescu, A., Boyer, E., Varanasi, K., Horaud, R.P., 2009. Surface feature detection and description with applications to mesh matching, in: Proc. of the IEEE Conf. on Comput. Vis. and Pattern Recognit.

Zhong, Y., 2009. Intrinsic shape signatures: A shape descriptor for 3d object recognition, in: Proc. ICCV Workshops, pp. 689–696.