Statistical criteria for shape fusion and selection

Alexandre Boulch   Renaud Marlet

Université Paris-Est, LIGM (UMR CNRS 8049), ENPC, F-77455 Marne-la-Vallée

http://sites.google.com/site/boulchalexandre   http://imagine.enpc.fr/˜marletr

Abstract—**Surface reconstruction from point clouds often relies on a primitive extraction step, that may be followed by a merging step because of a possible over-segmentation. We present two statistical criteria to decide whether or not two surfaces are to be considered as the same, and thus can be merged. They are based on the statistical tests of Kolmogorov-Smirnov and Mann-Whitney for comparing distributions. Moreover, computation time can be significantly cut down using a reduced sampling based on the Dvoretzky-Keifer-Wolfowitz inequality. The strength of our approach is that it relies in practice on a single intuitive parameter (homogeneous to a distance) and that it can be applied to any shape, including meshes, not just geometric primitives. It also enables the comparison of shapes of different kinds, providing a way to choose between different shape candidates. We show several applications of our method, experimenting geometric primitive (planeand cylinder) detection, selection and fusion, both on precise laser scans and noisy photogrammetric 3D data.**

I. INTRODUCTION

Numerous algorithms for surface reconstruction from point clouds have been developed over the past years, e.g., [2], [3], [4] just to cite a few. Many are based on the detection of simple geometric primitives: planes, but also cylinders, spheres, cones and tori. Algorithms for primitive extraction include RANSAC [1], region growing [2], [5] and Hough Transform [6], [7], [8]. All of these methods rely on a number of parameters, that are sometimes difficult to tune to fit the needs of the user.

This extraction step is most often crucial for the quality of the reconstruction and for the overall execution time. In many cases, there is actually a trade-off between reducing the number of segments to reduce computation time, and increasing the number of segments to enhance accuracy. To control over-segmentation, there often is an additional merging step that compares the geometric primitives underlying the segments.

Vosseleman et al. [9] merge planes if their equations are similar. It is equivalent to testing the similarity of plane normals and positions (w.r.t. some origin). It extends to any parametric surface but comparison between complex surfaces may be difficult due to possible high number of parameters, possibly with different dimensions (distances and angles). For instance, comparing cylinders requires tuning 5 parameters: 2 for cylinder direction, 2 for position and 1 for radius.

Addressing the parameter issue, Bughin proposes an a-contrario method to decide if a given point set can be interpreted as one or more planes [10]. However, it is specific to planes and defined for single range images.

A threshold on the Hausdorff distance can also be used to overcome the difficulty of parameterizing the comparison of complex shapes, as it relies on one parameter only. It has been used in particular for mesh comparison [11]. However the Haussdorf distance can be sensitive to outliers. Moreover, it

may requires computing many pointwise distances. Last, large sets of points are less likely to be considered as belonging to the same surface: the more pointwise distances are taken into account, the more likely outliers can affect the comparison test.

A better approach would be to use a criterion based on a non-parametric statistical test. Such tests have already been used, e.g., in the context of image segmentation [12], [13], [14]. We use it here for comparing shapes and fusing them.

In this paper, we present a new criterion for shape similarity relying on a single parameter (actually three, reduced to one in practice). It can be used both for shape fusion and shape selection. Besides, our method is less sensitive to outliers and also allows for a principled sampling strategy that greatly reduces the number of pointwise distances to compute. Comparison is not only faster; it also does not depend on the number of points in the segments for which surfaces are sought.

One original aspect of our method lies in the use of statistical tests, on a distance distribution, to decide whether two surfaces can be considered as equal. We also introduce a criterion inherited from robust statistics to reduce the number of points to test and consequently reduce the computation time and the dependence on the segment size.

The rest of the paper is organized as follows. Section II recalls basic statistical tests we rely on. Section III presents their application to shape comparison, fusion and selection. Section IV describes an efficient sampling strategy to reduce the size of the test. Section V explains how to meaningfully parameterize the test. Several applications and practical experiments are presented in Section VI: we merge over-segmented primitives via the identification of similar shapes in registered laser acquisitions; we illustrate shape selection on a scene with cylinders and planes, deciding whether a cylinder can be approximated by a plane; we show an example of primitive fusion on a noisy photogrammetric point cloud.

II. STATISTICAL TESTS FOR SURFACE COMPARISON

For deciding if two primitives detected in a point cloud can be merged compares the distribution of the distance of the support points to the primitives, we rely on statistical tests.

Let $X_1, \ldots, X_m$ and $Y_1, \ldots, Y_n$ be two random samples, i.e., two collections of i.i.d. random variables with respective distribution F_X and F_Y. Given observations $x_1, \ldots, x_m$ and $y_1, \ldots, y_n$, let $\hat{F}_X$ and $\hat{F}_Y$ be the corresponding empirical cumulative distribution functions.

$$F_X(t) = \mathbb{P}(X \leq t) \tag{1}$$

$$\hat{F}_X(t) = \frac{1}{m} \sum_{i=1}^{m} \mathbf{1}_{\{x_i \leq t\}} \tag{2}$$

The objective is to test the hypothesis $H_0: F_X = F_Y$.

(a) Segmentation with RANSAC-based method [1].　(b) Primitive types assigned to segments.　(c) Primitive type reassignment with our method.
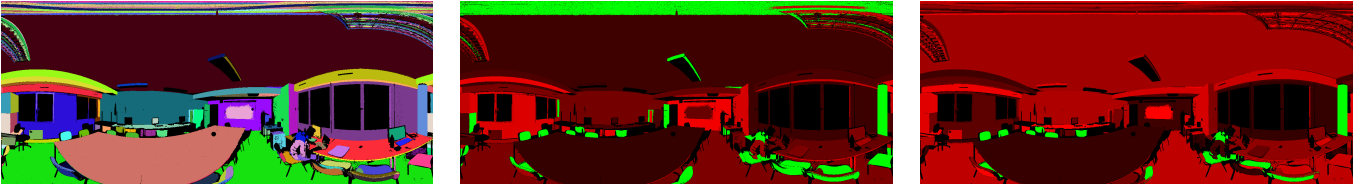
Fig. 1.　RANSAC-based shape detection (one color per segment) and assignment of primitive types. Shades of red correspond to planes, green to cylinders.

A. Kolmogorov-Smirnov test

The Kolmogorov-Smirnov test (KS-test) [15], [16] measures the deviation between two distributions. With this test, we mainly reject H_0 if the distributions have different shapes. For X, Y two samples of i.i.d. random variables, the test considers the statistic $KS_{stat} = D_{X,Y}$, which is the maximum deviation between the two empirical distribution functions:

$$D_{X,Y} = \sup_{t \in \mathbb{R}} |\hat{F}_X(t) - \hat{F}_Y(t)| \tag{3}$$

To compute $D_{X,Y}$, observations $x_1, \ldots, x_m$ and $y_1, \ldots, y_n$ are merged and sorted in increasing order into sequence $z_1, \ldots, z_{m+n}$. Let $S_k = \frac{1}{m}$ if z_k corresponds to an element x_i, and $S_k = -\frac{1}{n}$ if z_k corresponds to an element y_j. We have:

$$D_{X,Y} = \max_{1 \le k \le m+n} |\sum_{l=1}^{k} S_l| \tag{4}$$

According to the KS-test, we reject H_0 if

$$D_{X,Y} > c(\alpha) \sqrt{\frac{m+n}{mn}} \tag{5}$$

where $c(\alpha)$ can be tabulated according to the Kolmogorov-Smirnov distribution. Table I shows values of $c(\alpha)$ for common values of significance level α.

α	0.10	0.05	0.025	0.01	0.005	0.001
$c(\alpha)$	1.22	1.36	1.48	1.63	1.73	1.95

TABLE I.　$c(\alpha)$ FOR COMMON VALUES OF SIGNIFICANCE LEVEL α.

B. Mann-Whitney test

The Mann-Whitney test (MW-test) [17], [18] measures the interlacing of two distributions. This test mainly reject H_0 if the two distributions have different locations. The test is based on the MW-statistic $MW_{stat} = U_{Y,X}$ defined by:

$$U_{Y,X} = \sum_{\substack{1 \le i \le m \\ 1 \le j \le n}} \mathbf{1}_{\{Y_j > X_i\}} \tag{6}$$

If $U_{m,n}$ is a random variable of law $\mathcal{U}(m,n)$, which is the law of $U_{Y,X}$, and if $\min(m,n) \longrightarrow \infty$, then $\zeta_{m,n}$ defined by

$$\zeta_{m,n} = \frac{U_{m,n} - \frac{mn}{2}}{\sqrt{\frac{mn(m+n+1)}{12}}} \tag{7}$$

converges in law to the normal law $\mathcal{N}(0, 1)$.

For practical computation of the statistic, as in the KS-test, we merge and sort the values $x_1, \ldots, x_m$ and $y_1, \ldots, y_n$ into

a sequence $z = (z_1, \ldots, z_{m+n})$. If S_j is the rank of y_j in z, we compute the cumulated rank sum for Y:

$$R_Y = \sum_{j=1}^{n} S_j \tag{8}$$

$$U_{Y,X} = R_Y - \frac{n(n+1)}{2}. \tag{9}$$

For a test with significance level α, we use the quantile $1 - \frac{\alpha}{2}$.

III. SURFACE COMPARISON, FUSION AND SELECTION

A. Criterion for surface comparison and fusion

The above criteria can be used to test whether two surfaces extracted or defined from sets of points are likely to be identical. Let $\mathcal{S}_1$ and $\mathcal{S}_2$ be two surfaces, and let P_1 and P_2 be two sets of possibly noisy points associated respectively with $\mathcal{S}_1$ and $\mathcal{S}_2$. Let $d(p, \mathcal{S})$ be the distance from point p to surface $\mathcal{S}$. We consider the following two sets of values:

$$X = \{d(p_1, \mathcal{S}_1),\ p_1 \in P_1\} \cup \{d(p_2, \mathcal{S}_2),\ p_2 \in P_2\} \tag{10}$$
$$Y = \{d(p_1, \mathcal{S}_2),\ p_1 \in P_1\} \cup \{d(p_2, \mathcal{S}_1),\ p_2 \in P_2\} \tag{11}$$

If $\mathcal{S}_1$ and $\mathcal{S}_2$ are the same or almost the same surfaces, then the distribution of distances in X and Y are similar, and the above tests should not reject H_0, i.e., $F_X = F_Y$. Using a non-symetrized definition such as $X = \{d(p_1, \mathcal{S}_1),\ p_1 \in P_1\}$ and $Y = \{d(p_2, \mathcal{S}_2),\ p_2 \in P_2\}$ could lead to reject H_0 if the distribution of P_1 and P_2 are too different, even if $\mathcal{S}_1 = \mathcal{S}_2$.

In practice, P_1 and P_2 are generally known, and $\mathcal{S}_1$ and $\mathcal{S}_2$ are hypothesized, e.g., guessed using some kind of regression and/or RANSAC-like model detection. The statistical tests then provide a way to tell whether $\mathcal{S}_1$ and $\mathcal{S}_2$ are likely to be identical. If so, it means that P_1 and P_2 can be joined and associated to a single surface $\mathcal{S}_1$ or $\mathcal{S}_2$.

Provided we can hypothesize a common, merged surface $\mathcal{S}_3$ for $P_1 \cup P_2$, as well as corresponding point distances, this also provides a direct test for telling whether P_1 and P_2 support surface $\mathcal{S}_3$. For this, we replace Y in equation (11) by

$$Y = \{d(p_1, \mathcal{S}_3),\ p_1 \in P_1\} \cup \{d(p_2, \mathcal{S}_3),\ p_2 \in P_2\} \tag{12}$$

As an example, consider two planes $\mathcal{S}_1$ and $\mathcal{S}_2$, and the set of points P_1 and P_2 that was used to generate them. Considering $\mathcal{S}_3$, the regression plane of the set of points $P_1 \cup P_2$, the test tells whether it is legitimate to merge $\mathcal{S}_1$ and $\mathcal{S}_2$ into $\mathcal{S}_3$.

It is important to note that only a distance function to the surfaces is required here. Considering the problem of comparing and merging geometric primitives, it is easy to test planes, cylinders and other type of primitives. In fact, the two surfaces to compare can be of different kinds. The test actually also applies to mesh comparison and fusion.

B. Criterion for choosing surfaces or surface types

Sometimes it is difficult to assign a specific surface or surface type to a given set of points. Several candidates can be possible, with fitting scores that may not compare easily, or because a point set can be better fitted by one kind of surface than another. For instance, if a point set can be approximated by a plane, it can be approximated by a wide cylinder too and the cylinder is likely to have a smaller error.

Assuming there is a preference for certain directions or for certain kinds of surfaces (e.g., if planes are preferred to cylinders because they are considered simpler), then the surface comparison test also provides a way to choose between several candidates surfaces for a given set of support points.

For that, consider a set of points P and two surfaces $\mathcal{S}_1$ and $\mathcal{S}_2$ as candidate to approximate P. If the above test passes for $P_1 = P_2 = P$, then $\mathcal{S}_1$ and $\mathcal{S}_2$ can be considered as similar with respect to P. In this case, if additionally $\mathcal{S}_1$ is preferred to $\mathcal{S}_2$ for some other reason, it is safe to assign P to $\mathcal{S}_1$.

IV. EFFICIENT SAMPLING STRATEGY

Both tests have some dependence on the size of point sets: the distribution of distances to a surface is better estimated for a large number of points, but a small difference in the shape of the distribution then yields a rejection. Besides, the test computation time grows with the size of the point sets. These two issues can be addressed by an adequate point sampling.

Let $X = X_1, \ldots, X_n$ be a set of i.i.d. random variables with distribution function F, and $\hat{F}_n$ the empirical distribution function. Let $\beta \in]0,1[$ be a probability threshold expressing the confidence of the test for comparing $\hat{F}_n$ to F, and let $\epsilon > 0$ be a threshold for the deviation of the distributions. We want to estimate the minimum cardinal of X such that the maximum deviation between F and $\hat{F}_n$ is less than ϵ, with probability β:

$$\mathbb{P}(\sup_{t \in \mathbb{R}} |\hat{F}_n(t) - F(t)| \le \epsilon) \ge \beta \tag{13}$$

To estimate the number of points to use, we resort to the Dvoretzky-Kiefer-Wolfowitz (DKW) inequality [19]:

$$\mathbb{P}(\sup_{t \in \mathbb{R}} |\hat{F}_n(t) - F(t)| > \epsilon) \le 2e^{-2n\epsilon^2} \tag{14}$$

Using equation (13), we obtain a lower bound for n:

$$n \ge \frac{1}{2\epsilon^2} \ln(\frac{2}{1-\beta}) \tag{15}$$

In the following, we use $\beta = 1 - \alpha$, where α is the significance level of the test.

V. GIVING CONTROL TO THE USER

The method only depends on two parameters, a significance level α and a threshold ϵ for the distribution deviation. But the user has no control over the degree of similarity of the shapes. The slightest difference in the distributions leads to a rejection of the equality hypothesis. We want to introduce a parameter that is homogeneous to a distance and that defines some kind of acceptance level. Therefore, we consider a controlled noise that we add to the point set. The idea is that a larger noise yields a greater likelihood of similarity.
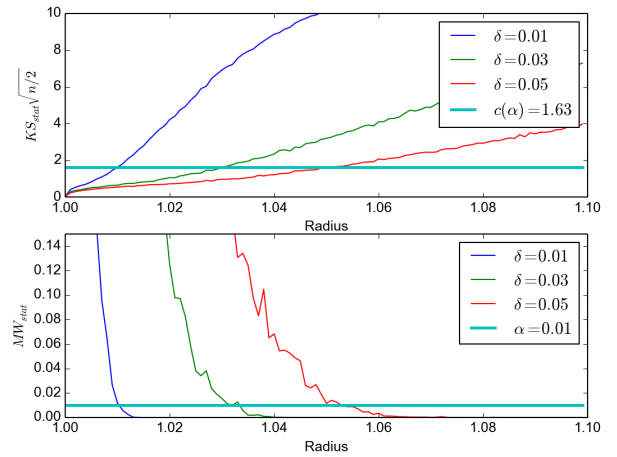

Fig. 2. Evolution of the KS and MW statistics for two co-centered spheres of different radii: the unit sphere and one with a variable radius in abscissa.

Given points sampled on two shapes, we empirically determine the standard deviation σ of the isotropic Gaussian noise to add to these points so that the shapes merge for the KS-test. We have experimented with four synthetic configurations: two parallel planes at a given distance h, two planes intersecting with a small angle $\theta \ll \pi/2$, two parallel cylinders with equal radius r and axis distance $h \ll r$, and two spheres with equal radius r and center distance $h \ll r$. ("$a \ll b$" indicates one order of magnitude or more, i.e., $10\,a < b$.)

Note that the situation is identical up to a scale factor. For parallel planes, σ thus depends linearly on h. Moreover, for cylinders and spheres it is enough to experiment with $r = 1$. Besides, given the above order of magnitude hypotheses, we experimentally observed that for intersecting planes, σ does not depend on θ. More importantly, for all configurations, σ depends linearly on the the mean distance δ between points sampled on a surface and the other surface (which is related to h), with a factor λ slightly depending on the configuration. It means, conversely, that given an average distance δ between points and shapes, adding noise on points with a standard deviation $\sigma > \lambda \delta$ is enough for the KS-test to predominantly decide merging. Experimentally, for $\alpha = 0.01$ and $\epsilon = 0.1$, we get $\lambda = 1.5$ for parallel planes and $\lambda = 1.6$ for the other configurations, with a standard deviation of 0.1.

In the following, we use δ rather than σ as the main parameter for deciding primitive fusion. We also rely on the value of λ for parallel planes, which provides a conservative lower bound: identical shapes may incorrectly not be merged, but different shapes are unlikely to be merged. From the point of view of the user, δ is the maximum distance (relative to scale) for which two parallel planes are allowed to merge.

VI. EXPERIMENTS

In this section, we present several experiments illustrating applications of our method. For all experiments, we use values $\alpha = 0.01$, $\epsilon = 0.1$, and thus $n_{\min} = 265$ points (cf. Eq. (15)).

A. Empirical comparison: KS-test vs MW-test

We empirically compare KS and MW tests on synthetic data. Figure 2 illustrates the response of the KS and MW

statistics for different values of δ when testing sphere similarity: we synthesize points on two co-centered spheres, one with unit radius, and the other one with a small radius variation. We consider 3 values for δ: 0.01, 0.03 and 0.05. The results are averaged over 100 drawings of point sets on the spheres for each value of δ. In this setting (with $\alpha = 0.01$ and $\epsilon = 0.1$), the KS-test seems to be more stable than the MW-test. Still, the two tests agree on the acceptance level for a given δ. In the following, we only use the KS-test to compare surfaces.

B. Merging regions grown from laser point clouds

Next we describe how our criterion can be practically used to extract consistent planes from a laser point cloud.

1) Algorithm: For segmentation into planar patches, we use a region growing algorithm similar to Bughin's [5], taking advantage of the image structure of laser acquisitions. A point is added to a growing region if it is at the border of the region and if it is less than $0.005\,\mathrm{m}$ far from the plane associated to the region. Point neighborhood is defined according to the 8-connectivity of the pixels in the depth image (spherical acquisition unfolded on a vertical cylinder, which is then unfolded on a rectangle). Seeds are selected randomly and only regions with at least 300 points are kept.

With this method, some planes can be over-segmented, i.e., if the regions are not connected in the image. Almost parallel planes can also be separate, e.g., posters on the wall or whiteboard. Both kinds of segments can be merged, in a controlled way.
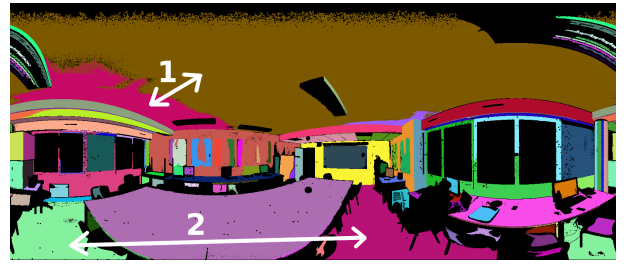
Primitive merging is performed greedily. First, the test is performed for all pairs of primitives. Those that pass the KS-test are sorted by increasing KS-statistic. (When using the MW-test, the order is reversed as the MW-statistic is higher for closer surfaces.) The pair of primitives with the lowest KS-statistic is merged and the sorted list of primitive pairs is updated with respect to the fusion. This process is repeated until no pair passes the test.

2) Evolution of fusion for different δ: Figure 3 illustrates the influence of parameter δ on the final segmentation. The scene is a laser acquisition of more than one million points. As expected, a growing δ leads to more fusions of primitives. Still, even with a big δ, the decision for merging primitives makes a lot of sense. Note that for $\delta = 0$ (not shown), no primitive happens to be merged due to the very high precision of the laser: any small variation in the shape of what we consider as planes is significant enough to reject fusion.

3) Primitive merging for separate acquisitions: To save memory and computation time, it is useful to treat laser scans in a hierarchical manner: first separately, possibly in parallel, then globally. Figure 4 illustrates the steps of this hierarchical segment-and-merge approach on three laser acquisitions. Regions are computed for each acquisition separately (4(a)), planes are merged inside each acquisition (4(b)), and finally primitives of all acquisitions are considered together for merging (4(c)).

C. Comparing apples and oranges

As already said, our test for surface similarity only requires the availability of distances of support points to the surfaces.

(a) Region growing result.

(b) $\delta = 0.01\,\mathrm{m}$.

(c) $\delta = 0.02\,\mathrm{m}$.

(d) $\delta = 0.03\,\mathrm{m}$.

Fig. 3. Evolution of plane fusion on a laser acquisition (1.1M points) for different values of δ. There is one color per segment. Examples of fusion between two consecutive values of δ are marked in white.

It applies to simple geometric primitives as well as arbitrary surfaces, including meshes. As explained in Section III-B, we can actually compare and choose between shapes that are different in nature, e.g., plane vs cylinder, as long as they are associated to point sets and corresponding distances.

To illustrate this application, Figure 1 shows the result of a RANSAC-based segmentation with both plane and cylinder detection [1]. The segmentation is greedy: once a primitive has been detected, its points are removed from the set of points to segment. In practice, as planes can approximate many surfaces, cylinders are generally sought first. However, an actual plane can then be misdetected as a flat fragment of cylinder. The question is, for each cylinder, whether the regression plane of its support points is a better primitive. Figure 1(a) shows the result of segment detection, and 1(b) the segment type (red

(a) After region growing in each acquisition separately.

(b) After primitive merging in each acquisition separately.

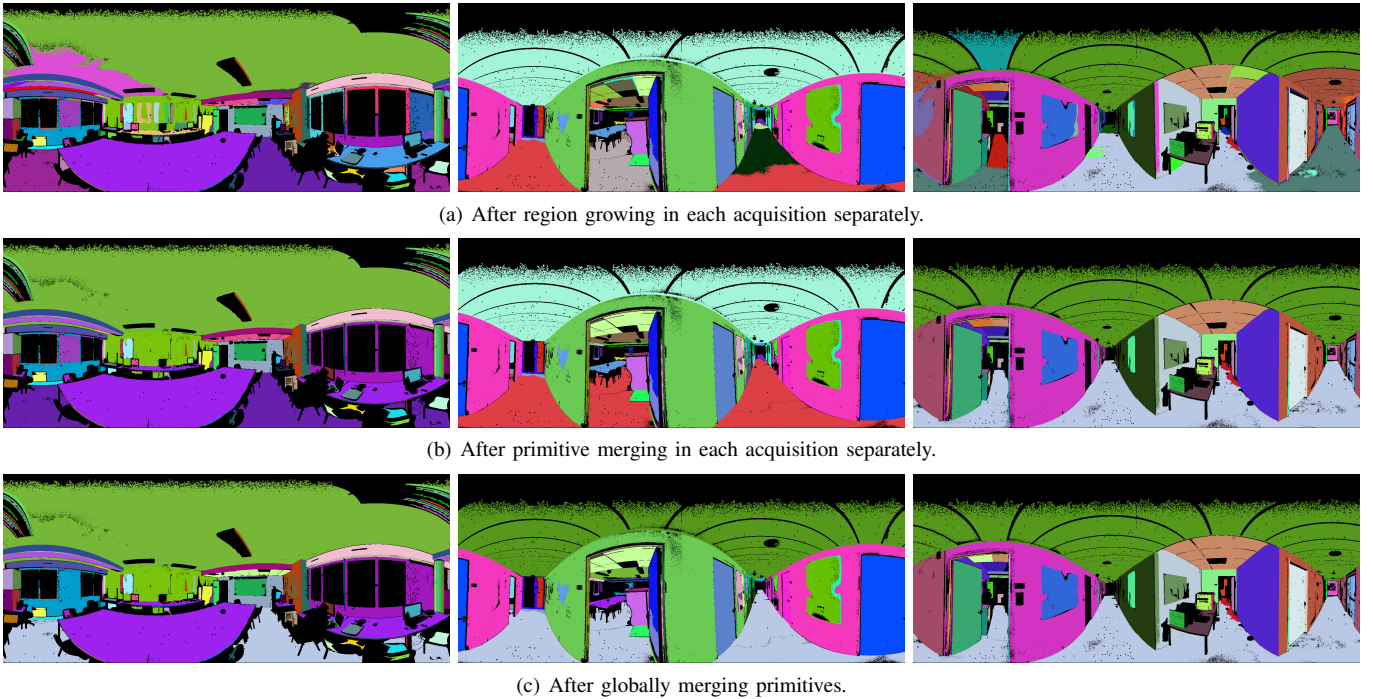(c) After globally merging primitives.

Fig. 4. Plane merging on multiple laser acquisitions (one color per segment): fusion is first applied to each acquisition separately, then globally (with $\delta = 0.01$).

for planes, green for cylinders). Note that some primitives are wrong. In particular, the door, part of the ceiling and the side of a large box are considered as cylinders rather than planes.

To address this issue, for each cylinder primitive, we compute its regression plane and compare the cylinder to the plane with respect to the (same) support points, with $\delta = 0.01$. According to the principle defined in Section III-B, if the test is positive, we choose the plane to best represent the segment. Figure 1(c) presents the result of such a choice. All segments that were erroneously considered as cylinders are now treated as planes. However, the chair backs and the pillar on the right are still correctly considered as cylinders.

D. Experimenting with photogrammetric data

We also experimented with photogrammetric 3D data, that are more noisy and whose density is less (locally) uniform than with laser scans. The input here is an actual point cloud, not (registered) depth images. We thus have to use a different approach than in Section VI-B to segment the acquisition into geometric primitives. However, we still resort to region growing. A point is added into a segment if its normal and the normal of its neighbors are nearly the same. Normals are computed using the approach proposed by Boulch and Marlet [20], which is robust to sharp angles. This results in segments with smooth curvature that do not necessarily correspond to simple geometric primitives. Then, for each segment, we compute its regression plane and, possibly, a cylinder. The cylinder is sought in a RANSAC fashion, picking random pairs of points, hypothesizing a corresponding cylinder (based on the normal of the two points), checking how many points fit the cylinder, and keeping the cylinder with the most points. After that, we use the principle defined in Section III-B to choose between the plane and the cylinder. Finally, cylinders and planes are merged separately, as defined in Section VI-B1.

This process is illustrated on Figure 5 using 3D data of the *Pavillon de l'Aurore* in Sceaux, reconstructed automatically by photogrammetry (from pictures) with the method of Moulon et al. [21], yielding a 3D point cloud with 1.6M points. As can be seen, the four segment of the central part are properly detected as cylinder fragments, and properly merged.

VII. CONCLUSION

We have presented two variants of a general statistical test to assess the similarity of surfaces underlying point clouds. We have explained how it can be used to compare and merge surfaces, in particular geometric primitives, and also to choose between different shape candidates. We have shown how a small number of points can be enough the evaluate the test, also providing insensitivity to point set size, and we have introduced a meaningful parameter δ for the user to tune the acceptance level of the test. The strength of our approach lies in its single (in practice) and intuitive parameter, as well as in the generality of the test because it applies to any surface as long as we can compute distances. We have extensively illustrated how our criterion could be used to consistently segment a point cloud into primitives, possibly of different kinds, using either clean and precise laser scans as well as noisy photogrammetric data.

In these experiments, the merging strategy is driven by test statistics but it is still iterative and greedy. It is future work to use our criterion in a more principled algorithm for primitive merging, with better optimality guarantees. Application to other shape regularities can also be though of, such as polygonalization of segment borders.

REFERENCES

[1] R. Schnabel, R. Wahl, and R. Klein, "Efficient RANSAC for point-cloud shape detection," in *Comp. Graph. Forum*, vol. 26, no. 2, 2007, pp. 214–226.
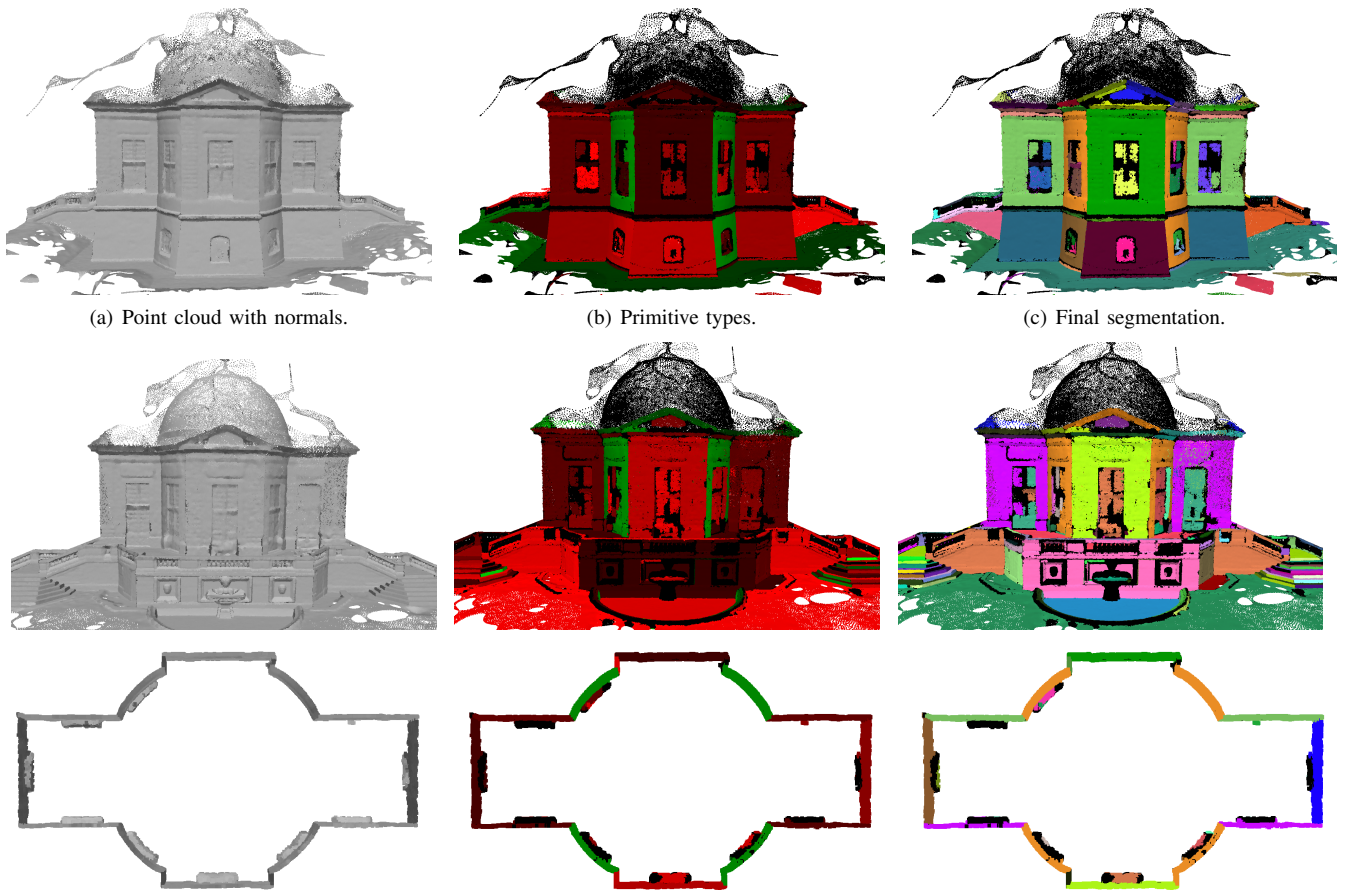
Fig. 5. Photogrammetric data, geometric typing of segments (red for planes, green for cylinders), and segment fusion (different colors for different primitives).

[2] A.-L. Chauve, P. Labatut, and J.-P. Pons, "Robust piecewise-planar 3D reconstruction and completion from large-scale unstructured point data," in *CVPR*, 2010, pp. 1261–1268.

[3] Y. Li, X. Wu, Y. Chrysathou, A. Sharf, D. Cohen-Or, and N. J. Mitra, "Globfit: consistently fitting primitives by discovering global relations," in *ACM Trans. on Graphics (TOG)*, vol. 30, no. 4, 2011, pp. 52:1–52:12.

[4] F. Lafarge and P. Alliez, "Surface reconstruction through point set structuring," in *Proc. of Eurographics*, Girona, Spain, 2013.

[5] E. Bughin, A. Almansa, R. Grompone von Gioi, and Y. Tendero, "Fast plane detection in disparity maps," in *ICIP*, 2010, pp. 2961–2964.

[6] P. V. C. Hough, "Method and means for recognizing complex patterns," *U.S. Patent*, no. 3.069.654, 1962.

[7] J. Knopp, M. Prasad, G. Willems, R. Timofte, and L. J. V. Gool, "Hough transform and 3D SURF for robust three dimensional classification," in *ECCV*, ser. LNCS, vol. 6316, 2010, pp. 589–602.

[8] M.-T. Pham, O. J. Woodford, F. Perbet, A. Maki, B. Stenger, and R. Cipolla, "A new distance for scale-invariant 3D shape recognition and registration," in *ICCV*, 2011, pp. 145–152.

[9] G. Vosselman, B. G. Gorte, G. Sithole, and T. Rabbani, "Recognising structure in laser scanner point clouds," *International Archives of Photogrammetry, Remote Sensing and Spatial Information Sciences*, vol. 46, no. 8, pp. 33–38, 2004.

[10] E. Bughin, "Towards automated, precise and validated vectorisation of disparity maps in urban satellites stereoscopy," Ph.D. dissertation, ENS Cachan, 2011.

[11] N. Aspert, D. Santa-Cruz, and T. Ebrahimi, "Mesh: measuring errors between surfaces using the Hausdorff distance," in *IEEE Int'l Conf. on Multimedia and Expo (ICME 2002)*, vol. 1, 2002, pp. 705–708.

[12] L. Wolf, X. Huang, I. Martin, and D. N. Metaxas, "Patch-based texture edges and segmentation," in *ECCV*, 2006, LNCS 3952, pp. 481–493.

[13] C. Fiorio and R. Nock, "A concentration-based adaptive approach to region merging of optimal time and space complexities," in *BMVC*, 2000, pp. 1–10.

[14] R. Nock and F. Nielsen, "Semi-supervised statistical region refinement for color image segmentation," *Pattern Recognition*, vol. 38, no. 6, pp. 835–846, 2005.

[15] A. N. Kolmogorov, "Sulla determinazione empirica di una legge di distribuzione," *Giornale dell'Istituto Italiano degli Attuari*, vol. 4, no. 1, pp. 83–91, 1933.

[16] N. Smirnov, "Table for estimating the goodness of fit of empirical distributions," *The Annals of Mathematical Statistics*, vol. 19, no. 2, pp. 279–281, 1948.

[17] H. B. Mann and D. R. Whitney, "On a test of whether one of two random variables is stochastically larger than the other," *The Annals of Mathematical Statistics*, vol. 18, no. 1, pp. 50–60, 1947.

[18] F. Wilcoxon, "Individual comparisons by ranking methods," *Biometrics bulletin*, vol. 1, no. 6, pp. 80–83, 1945.

[19] A. Dvoretzky, J. Kiefer, and J. Wolfowitz, "Asymptotic minimax character of the sample distribution function and of the classical multinomial estimator," *The Annals of Mathematical Statistics*, pp. 642–669, 1956.

[20] A. Boulch and R. Marlet, "Fast and robust normal estimation for point clouds with sharp features," *Comp. Graph. Forum*, vol. 31, no. 5, pp. 1765–1774, Aug. 2012.

[21] P. Moulon, P. Monasse, and R. Marlet, "Global fusion of relative motions for robust, accurate and scalable structure from motion," *IEEE International Conference on Computer Vision (ICCV)*, Dec. 2013.