EXTRACTION OF PLANAR PATCHES FROM POINT CLOUDS TO RETRIEVE DIP AND DIP DIRECTION OF ROCK DISCONTINUITIES

R. Roncella, G. Forlani
Dept. of Civil Engineering, Parma University, Parma, 43100, Italy
rroncell@nemo.unipr.it, gianfranco.forlani@unipr.it

KEY WORDS: Risk assessment, Geomorphology, Photogrammetry, Automation, Modeling, DEM/DTM, Segmentation, RANSAC, Multi-Resolution.

ABSTRACT:

Laser scanners and image matching algorithms produce large amounts of 3D data, allowing detailed and accurate descriptions of objects. This is true, though, only if interactive or automated software tools are available, to extract the information relevant to the specific task from raw 3D data sets without too much user interaction. In the last two years our group has been involved in applying new surveying methods to determine the dip and dip directions of discontinuities in rock faces; knowledge of these parameters, that can be computed from an high resolution DSM of the rock, allows for a stability analysis to be carried out. In this paper we describe an automated approach to extract from a dense DSM of a rock face, obtained by photogrammetry or by laser scanning, the discontinuity planes of the rock and compute their parameters. A multi-resolution DSM pyramid is generated from the original DSM; in a hierarchical scheme, the DSM is segmented in planes at each level of the pyramid, using RANSAC with topological constraints. Two clustering steps are also executed, to balance the fragmentation of the planes and the accuracy of the surface approximation. Finally, dip and dip directions are computed from the normal vector to each plane. The proposed method, applied to a test site, gave results equivalent to those of a traditional survey.

1. INTRODUCTION

A comprehensive stability analysis of rock slopes is essential to foresee or at least evaluate the likelihood of rock avalanches. Demand for fast and efficient data acquisition techniques enabling the application of statistical methods in the stability analysis is raising with the improvement of numerical modelling and the increasing number of sites where the risk must be assessed (i.e. mountain roads and villages, skiing facilities, old and active quarries, etc).

The rock mass structure and the distribution of the discontinuities are the main parameters affecting rock stability. Discontinuities are surfaces of weakness within the rock, from a micro scale up to a macro scale. The geo-structural rock characterization can be obtained by surveying the distribution of the discontinuities orientation and their spacing over the rock surface, extending their pattern into the mass volume by statistical techniques.

The orientation of a discontinuity plane is defined by two angles: dip and dip direction. Dip is the angle between a horizontal plane and the discontinuity plane; dip direction is the azimuth of the projection of the gradient vector of the discontinuity plane on the horizontal plane (see Figure 1). Dip and dip direction are measured on site by a geological survey, carried out using compass along scan lines, recording the location of the measurement point on sketches or images.

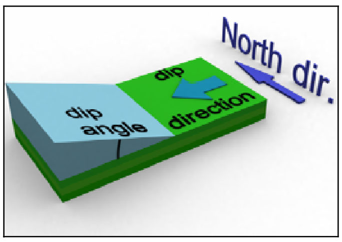

Figure 1: Dip and dip direction of a discontinuity plane.

Alternative methods based on photogrammetry have been proposed (Harrison et al, 2000) to overcome the limits of the on-site surveys. A discontinuity plane is identified by geometric data only: they can be obtained as well by means of simple computations, should the equation of the discontinuity plane known. If the North, East, Elevation coordinates of a number of well distributed points on the discontinuity plane are determined, its location (e.g. the gravity centre) can be computed; fitting a plane will return the dip and dip direction of the discontinuity from the components of the vector normal to the plane.

Applying this procedure to a large number of rock discontinuities, a spatial data base of their location and orientation would be generated. Besides, their spacing and density may be evaluated by clustering planes with similar orientation and analysing their spatial distribution.

To be a real improvement with respect to the geological survey, an alternative system based on 3D data only should dispense with accessing the rock; dip and dip direction should be recovered on all significant discontinuity in the rock, with an accuracy at least matching that of the compass. Overall costs should be smaller than those of a comparable traditional survey. These requirements can be met either by photogrammetry or by terrestrial laser scanning; a detailed analysis of pro and cons of both techniques is beyond the scope of this paper. In the following, we concentrate rather on the automation of the extraction of discontinuity planes, which applies to a generic point cloud and can therefore be used independently of the data acquisition technique.

2. A PROCEDURE FOR THE MEASUREMENT OF DIP AND DIP DIRECTION

Under the INTERREG III project ROCKSLIDETECT, in cooperation with the rock mechanics group of our Department (Ferrero et al, 2004), we set up the development of a semi-automated photogrammetric system for the retrieval of rock discontinuities. The goal is that a geologist, with basic

surveying knowledge, using a good resolution digital camera and following specific image acquisition guidelines should acquire and process the data with a package of software programs, retrieving the distribution of discontinuities on the rock. The software must provide automated or semi-automated image orientation, DSM generation, DSM segmentation (discontinuity planes extraction) and finally dip and dip direction computation.
Since the last two steps are actually independent of how the DSM is produced, the system will straightforwardly accommodate any change in the technology used for data acquisition.
To verify the suitability of the method, a test site has been set up on the South West face of Corma di Machaby (Arnad, Italy). A rock section stretching for about 150 meters in width and 90 meters in height (Figure 2) was surveyed with geological compass to provide reference data.

Figure 2: The Corma of Machaby test site.

The feasibility of using photogrammetry to retrieve dip and dip direction has been shown in (Roncella et al, 2005) using either manual restitution and image correlation. As far as the automatic procedure is concerned, in the same paper it was shown that S&M (structure and motion) reconstruction and dense point matching can provide a detailed DSM, on a par with a laser scanner. A DSM segmentation scheme, to automatically retrieve discontinuity planes was also presented, although early results did not match very well the reference data. In the following, we present a semi-automatic and a fully automatic procedure for the extraction of discontinuities from the point cloud, successfully applied to the Machaby test site. Results refer to a DSM generated by image correlation using the software Virtuozo by Supresoft, from a stereo pair manually oriented on g.c.p. Images were taken with a Nikon D100 at an image scale of about 1:3300.

3. SEMI-AUTOMATIC DETERMINATION OF THE DISCONTINUITY PARAMETERS

Once the DSM is available, we have to extract planar surfaces from it, classifying each in terms of goodness of fit (mean square error), cardinality, location of the gravity centre and finally dip and dip direction.
A simple software tool was created to help a geologist with the manual extraction of discontinuities. Using a previously oriented image, the point cloud is back-projected onto the image frame. The user draws a polyline on the image, bordering an area enclosing discontinuities; the algorithm searches in object space the point within the polygon which is closest to the camera centre; the 3D points of the polygon within some threshold of the closest point in the viewing direction are selected.
To determine the dip and dip direction angles of the selected portion of the rock façade, a RANSAC-based algorithm (see next section) extracts all planes in the region.
Dividing the rock façade into different areas with roughly the same orientation and letting the algorithm to execute a detailed segmentation, results very close to the reference data were obtained, without significant computational effort.

4. AUTOMATIC DISCONTINUITY EXTRACTION

Even if the semi-automated procedure leads to reliable results, a fully automatic procedure is desirable, to sample the rock face in a regular way, producing a significant statistics of the main families of discontinuities. The interactive method may be used as a necessary check in some areas, to validate the results.
Due to the amount of points, which may be in the order of several hundred thousand or even millions, just trying to use the basic RANSAC procedure of the semi-automatic approach would be computationally unfeasible. Therefore, a multi-resolution scheme was developed, where first a DSM pyramid is generated, then each level is processed and its results transferred to the next.

4.1 DSM pyramid generation

The DSM pyramid levels are generated in a voxel approach, enclosing the point cloud in a box and dividing it in cells of different size at each level. The pyramid generation starts from the bottom level, where the cells have the smallest size: the points falling in each cell (we choose the cell size to get 4-5 pts/cell on average) are substituted for by their gravity centre. From one level to the higher, 8 cells (or more if a higher order of decimation is required) of the lower level are grouped into a new cell; the point associated to the new cell is the gravity centre of the 8 parent cells. With a reduction factor of four at each level, the DSM segmentation becomes soon computationally feasible with RANSAC. Because of the way the 8 cells are grouped, the smoothing is isotropic; the average distance of the gravity centres from the original surface at each level remains almost unchanged along the pyramid. As can be seen from Figure 3, the differences of the original point cloud with respect to the lowest level of the pyramid.

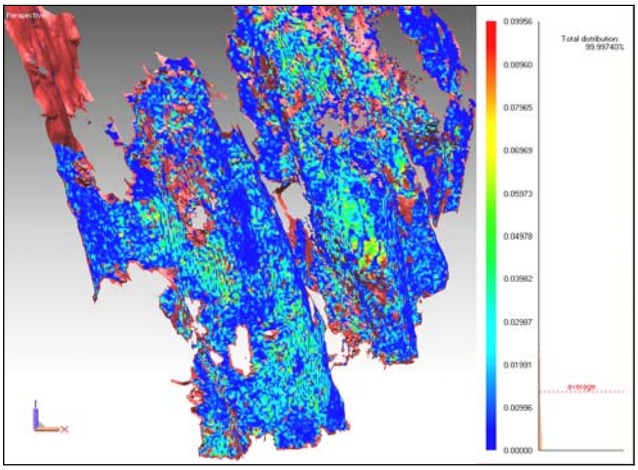

Figure 3: Distances between the full resolution DSM and the DSM at the highest pyramid level; mean = 1.5 cm.

A smoothed version of the original surface is generated in the higher levels of the pyramid. This helps the extraction of truly planar macro areas of the rock surface (i.e. the "best consensus sets" will be planes tangent to the surface rather than cutting it), avoiding a too complex segmentation of the surface which can be deceptive in the next steps. Moving to the lowest level of the pyramid, the main structure of the facade morphology is therefore already recovered and will act as a constraint on the subsequent segmentations.

4.2 Planes extraction and fitting.

From the top to the bottom of the pyramid, plane surfaces are extracted by RANSAC (Fischler and Bolles, 1985) at every level. A minimum subset of points (three) is randomly sampled from the point cloud, defining a plane; the coefficients of the plane equation in homogeneous form are computed by SVD:

$$ax + by + cz + d = 0 \qquad (\text{subject to } \|(a,b,c,d)\| = 1) \tag{2}$$

Counting the number of points whose distance from the selected plane is less than a threshold, a score is computed. Randomly selecting different minimal subsets (different planes), after a number of trials the plane with the best score is chosen. Removing from the point cloud the points within the threshold of the best plane, a new RANSAC step is executed and a new plane extracted. Ideally, this would end in selecting all planes, from the largest to the smallest, until no points remain to assign or no planes with significant consensus can be found.

The number of iterations required to ensure that at least one of the random subset is free from outliers depends on the probability to get at least a minimum set free from outliers (often set close to 100%), the size of the subset (i.e. 3 points), the expected outliers percentage.

In the context of this application, a minimal subset free from outliers would ideally be any tern of points actually chosen on a true plane patch of the rock surface. In most cases, on the contrary, randomly selecting points will define planes cutting through the rock surface, rather than being tangent to it. Obviously, points within the threshold are in the former case truly "good ones" to represent discontinuities, in the latter they are "bad ones". In RANSAC terminology, though, they are both inliers: therefore also planes "cutting through" with a large number of inliers may be selected as "best consensus" planes. The outliers percentage, set with an adaptive strategy, is updated at every plane selection; it therefore decreases within a level and is reset to 99% moving to the next one.

Using a relatively large threshold, the smoothness of the surface allows the identification of the main features (macro planes) of the rock structure. In the next steps of the procedure, only single macro areas are passed to RANSAC, so that a narrower threshold can be used since a small range of plane orientations is expected.

In the current implementation, once a "best consensus" plane is chosen, its equation is estimated in a l.s. adjustment with all points within the threshold. Due to the different plane orientation and position after adjustment, a new set of inliers is computed. Adding and removing observations to the system at each iteration, until no new points get in leads to a stable solution (i.e. to the identification of the same planes independently of the randomly selected seed points).

4.3 Topological constraints

After the first tests, it became clear that for the described approach to succeed, the distance threshold should be complemented by other topological and geometric conditions; otherwise, especially at the pyramid top, where the threshold is large, RANSAC may still select planes cutting through the surface because they have large consensus. Different criteria were implemented in the algorithm to avoid this:

- **connectivity**: with a kind of region growing from a seed point, a subset is created where every point must be closer than a threshold to its nearest neighbour in the set; this should prevent points on rock structures far apart from each other to be joined in the same plane;
- **convexity**: when the convex hull area differs significantly from that of the inliers footprint on the selected plane, the inliers region has too many holes or branches, i.e. it is likely to be a cutting plane rather than a plane tangent to the rock surface;
- **shape**: if the ratio between the principal moments of inertia is larger than a threshold, the shape of the inliers region is too elongated and the estimation of the plane parameters may be too uncertain.

4.4 Clustering based on orientation

Although sound, applying all these criteria tend to rule out too many selected planes, leading to an undesirable piecewise subdivision of the surface. Since at next pyramid level the previously determined planes may be further subdivided, the technique can lead to a final result with thousands of different planes. This outcome is not optimal for the geo-structural analysis, mainly because a big discontinuity family (i.e. one with a large area) would have a lot of different planes with almost the same orientation; this would overshadow the existence of a smaller but still remarkable family. To avoid this, at every iteration step all planes extracted are compared with the contiguous ones, to test whether their orientation is about the same. If so, their points are fused and a new plane is estimated. Moreover, if some points become outliers with respect to the new plane, they are temporarily assigned to the closest plane and further investigated in the next level.

To improve the algorithm performance, a hierarchical clustering has been applied (Anderberg, 1973). The method starts with each item as a separate cluster (the leafs of the cluster tree); under a given metric, the mutual distances between the objects are stored in the matrix D. A cluster tree, which can be visualized by a dendrogram, is built with a linkage function, which measures the between-clusters distances. First, the closest pair of objects is fused in a new cluster; then, the procedure is repeated, fusing each time the two closest clusters, until only one cluster is left, at the top of the hierarchical tree.

To evaluate the between-clusters distance, different principles may be applied. We used the shortest distance, where the linkage function is defined as:

$$L(r,s) = min\left(dist\left(p_{ri}, p_{sj}\right)\right) \tag{3}$$

where p_{ri}, p_{sj} are respectively the i^{th} object in cluster r and the j^{th} object in cluster s and *dist* is the distance in the dip-dip direction plane. Since at each step two cluster are joined, if there are m objects, after m steps (levels) the cluster tree is completed.

The cluster tree provide a sort of multi-resolution picture of the data set: traversing the dendrogram (cutting the tree branches) at a different level will originate different clusters (Figure 4 shows an example of hierarchical clustering).

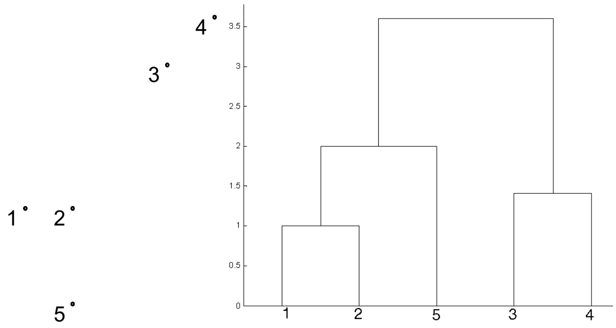

Figure 4: a) a set of 5 objects; b) cluster tree representation.

The optimal level for this cut may be inferred by comparing the length of each link between cluster pairs in the cluster tree with those of neighbouring links: if this is quite the same, the data clustered at this level show great similarity (i.e. an high level of data consistency), otherwise the link appear to be inconsistent. To find the appropriate level, an inconsistency coefficient can be computed for each link (*r, s*) between cluster *r* and *s*, as:

$$Y(r,s) = \frac{\left(L(r,s) - \overline{M}\right)}{S} \tag{4}$$

where L(*r, s*) is the linkage distance from equation (3) and M and S are respectively the mean and the standard deviation of the length of the links included in the calculation. The greater its value, the less the clusters connected by the link are similar to the others.

Since we look in a group of neighbouring planes for those with similar orientation, every extracted plane is characterized by its normal (defined by dip and dip direction). The mutual distances on the dip-dip direction plane are computed for every surface pair. Next every pair of objects are linked together using a proximity criteria (e.g. planes with smaller distance in the dip-dip direction plane are grouped).

Using different thresholds for the inconsistency coefficient leads to different clusters: at different scales, the surfaces are grouped upon the parallelism of their normals. The threshold can be chosen a priori by the user (i.e. only planes with normals parallel up to a certain value are grouped together) or can be adaptively established using information from the cluster tree.

Finally, it's possible to evaluate the goodness of the clustering with the cophenetic correlation coefficient, defined as:

$$c = \frac{\sum_{i<j}\left(D_{ij} - \overline{D}\right)\cdot\left(L_{ij} - \overline{L}\right)}{\sqrt{\sum_{i<j}\left(D_{ij} - \overline{D}\right)^2 \cdot \sum_{i<j}\left(L_{ij} - \overline{L}\right)^2}} \tag{5}$$

where:

- D*ij* is the distance between the original objects *i* and *j*;
- L*ij* is the link length (the linkage function value) between a pair of clusters containing the objects *i* and *j* when they were first joined in the cluster tree;
- D, L the mean values of D*ij* and L*ij*.

The cophenetic coefficient measures the distortion of the suggested classification, indicating how readily the data fits into the clustering structure: the closer the magnitude of its value to 1, the higher the quality of clustering structure.

Since the method consider only the relative spatial distribution of the points, the same inconsistency factor and cophenetic coefficient would be obtained should the points locations modified by a scale transformation. This means that, in terms of cluster quality, it is irrelevant whether the minimum distance between points (the planes) or between clusters is, in absolute terms, larger than a threshold: planes with rather different orientations (i.e. representing different discontinuity families) may end up in a single cluster. To avoid this, in the computation of the D matrix the distances above some user defined threshold are appropriately increased: this enhances the separation between clusters.

4.5 Splitting up

Clustering of neighbouring planes based on similarity of the orientation is necessary to avoid the segmentation of the DSM to end up in a messy proliferation of small patches. Sometimes, though, clustering based on adjacency and orientation produces some fuzzy-shaped regions. Take as an example the sketch of Figure 5a, where the pink U-shaped region may be thought as being generated by a plane cutting through the rock; although geometrically consistent (the region actually fits a plane), this may lead to an incorrect interpretation of the discontinuity from a geostructural point of view (e.g. the green edge may indeed separate two discontinuity planes) and must therefore be split.

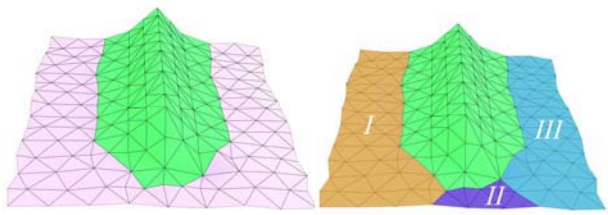

Figure 5: a) U-shaped region (incorrect aggregation during hierarchical clustering); b) Splitting region by k-means clustering.

This problem is dealt with in another data processing step, where each aggregated plane may be split by k-means clustering, trying to generate compact (convex) regions, each possibly corresponding to a single discontinuity plane, ideally ending up with something like the new subdivision shown in Figure 5b.

The k-means algorithm (Hartigan, 1975) divides the observations of a data set into mutually exclusive clusters, optimising some objective function. Unlike the previous approach, a single level of clusters is created, where the observation values and not only their proximities are considered. The goal is to divide the obiects (i.e. the 3D points lying on the selected plane) into K clusters such that some metric relative to the centroids is minimized: this shold break elongated regions, cutting links between adjacent regions are broken.

K-means uses an iterative method where, at every iteration, the algorithm moves the observations between clusters until the sum of the distances of each point from the cluster centroid cannot be further decreased. As a matter of fact, k-means achieve the most compact form of the clusters trying to limit the number of subdivisions of the points analysed.

The clustering subdivision (i.e. how many clusters must be used in the splitting phase) can be defined by the user or adaptively determined by the algorithm: in the latter case different partitions are tried and the best is chosen; even if computationally more demanding, this provides the best results.

4.6 Parameter's setup

The procedure described above requires many different parameters to be specified by the user. As far as the pyramid generation is concerned, the smallest size of the voxels and the lowest resolution accepted (and so how many pyramid levels) depend upon the computational effort which can be considered acceptable and on the minimum number of points required to describe completely the surface; these parameters are easily set after some tests.
The RANSAC algorithm requires two parameters and some tests to obtain the right tuning. In principle, the correct outlier percentage in the first level (i.e. when macro areas are extracted) depends on how many macro-planes there are; with the adaptive evaluation of the outlier percentage at first a small inlier percentage (close to 0%) is assumed; at every iteration the number of inlier is estimated and the outlier percentage updated. If the iterations performed so far are larger than those required with that outlier percentage, the algorithm stops, otherwise iterations continue. A more critical point is the correct threshold for the distances to the randomly extracted plane at every level of the pyramid: large values mean less computations and larger regions, but important details may be missing. On the other hand, using a too small threshold tends to obtain a over divided surfaces, as already pointed out.
A series of comparative tests has been performed to find out whether there exist an optimal threshold value and if different rock morphology requires different thresholds. Summarizing, even if the rock surface geometry varies (some area may be smoother, some other may be rough) a common threshold value was found acceptable. Indeed, from a geo-structural standpoint, on smooth surfaces the information is captured by very few planes, while in rough areas, with sharp edges, it is important to pick every plane. A single threshold, not too narrow but not too large (see below) will do just that.
As far as extracting discontinuities is concerned, the used threshold matters only in the zero level of the pyramid; here it has the greatest importance, since it determines the actual planes extracted: from the tests carried out and relying on the opinions of some geologist, a value between 10 and 15 cm will track all major discontinuity planes in most cases.
Also the clustering steps require correct understanding of the different parameters involved: a too high aggregation in the hierarchical tree could bring poor results, while allowing a too dramatic splitting by the k-means clustering may end up with an extremely divided, hard-to-analyse surface as well as requiring a greater computational effort. Currently a good balance of the two algorithm has been achieved after some tests, but in the future a less empirical approach should be developed: to this aim, we plan to use fuzzy logic.

4.7 Algorithm flow-chart and summary of the procedure

Once all planes in a level have been computed, clustered and split to obtain a more consistent arrangement, they are checked at the next resolution level. The threshold for the acceptance of inliers is therefore made smaller, to refine the local approximation of the planar regions. At the same time, when all planes of the *i-1* level have been analyzed at level *i*, a new hierarchical clustering step is performed, mainly to join again the regions at the boundary of a plane that were incorrectly subdivided in the *i-1* level of resolution. Figure 6 summarizes the flow-chart of the processing steps.

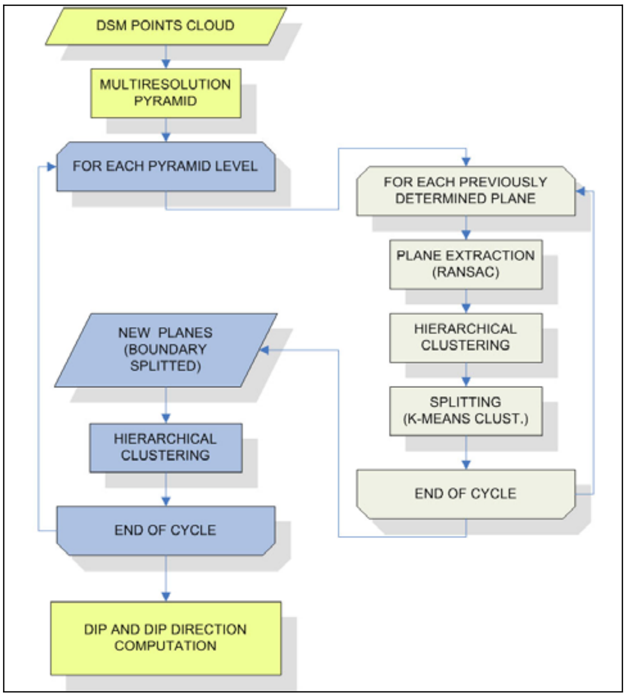

Figure 6: The flowchart of the plane extraction algorithm.

4.8 Results evaluation and comparisons

The final DSM segmentation of the Machaby test site consisted of 43 different planes, selected with a threshold in the maximum resolution step of 15 cm. Running the planar segmentation, implemented in Matlab, took about 30 minutes on a Athlon 2 GHz workstation on a 1.3 M points DSM. Being the software not yet optimized, significant improvements in computing time are expected.
Figure 7 shows the final segmented DSM: even if some clustered plane seems not correctly determined, the geostatistical analysis points out that the main discontinuity families are correctly drawn.
Figure 8 shows the stereoplot of the discontinuity families (top) and the associated main orientations (bottom): on the left size the results of the on-site survey by the geologist, on the right side the result of the automatic segmentation. The stereoplot shows the contour lines of the dip-dip direction distribution in the plane azimuth - elevation on the horizon: vertical planes are close to the border of the circle. The orientation plot (top) represents the intersection of the normal to the discontinuity plane with the unit sphere: here, vertical planes cross the circle close to the pole of sphere. The similarity of both graphs, according to geologists, is very good. Only a small family of discontinuities is missing, which is sub-horizontal: this is most likely because the images used for DSM generation were taken from ground only, with viewing directions almost co-planar. Therefore, any horizontal plane would likely be occluded or viewed under an unfavourable angle.
Besides, a discontinuity family (the black line from south to north in the center of the bottom-left diagram) not highlighted by the the geological survey was captured by the automatic segmentation procedure.

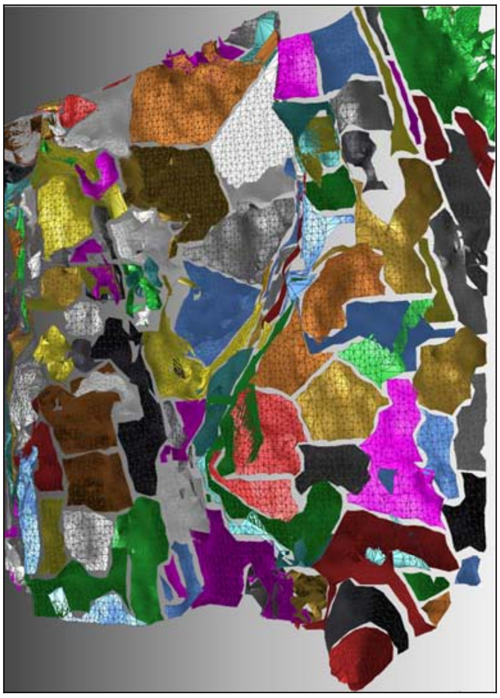

Figure 7: The final segmentation of the DSM: 43 different planes were selected.

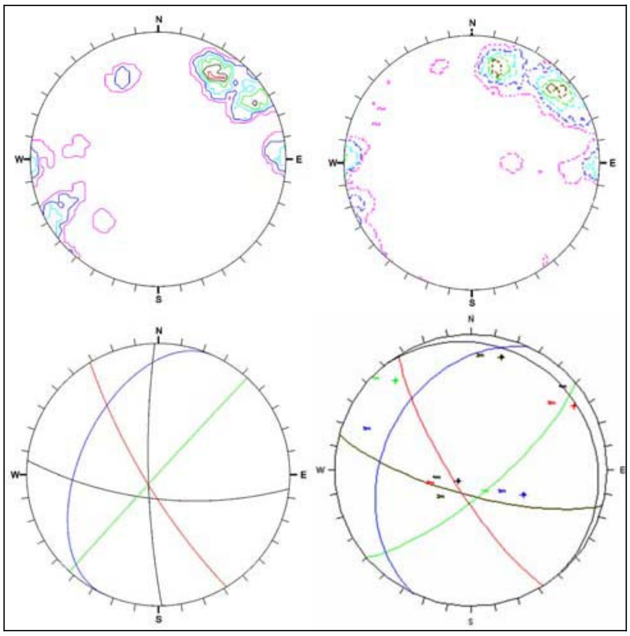

Figure 8: The stereoplots of the discontinuity planes. Left: the outcome of the automatic DSM segmentation; right: results of the manual geological survey.

5. CONCLUSIONS AND PERSPECTIVES

An highly automated technique to extract from the DSM of a rock face its discontinuity planes has been presented. It relies on a multi-resolution analysis of the DSM pyramid; at each level, planes are identified in the point cloud with a RANSAC-based algorithm; since too much fragmentation may arise, which is not optimal for the classification of discontinuities with geostructural statistics, planes are first merged at the proper threshold in a cluster tree, then divided again whenever clustering based on proximity and parallelism led to the aggregation of planes in regions with shapes misleading for geostructural analysy.

The procedure has been successfully tested on the Machaby site, were a large set of reference data was available.

Although this is encouraging, some remarks are necessary. The variety of rock morphology in the sites where the technique is supposed to be applied suggests that many other tests are necessary before we can be confident of the robustness of the procedure. As far as the true automation level of the procedure is concerned, some processing parameters are actually set on visual inspection of the results. A better understanding of the characteristics of an ideal segmentation, drawn from discussion with the geologists, to translate their expertise in more strict and clear constraints for the segmentation may help not only to rubustify it, but perhaps will suggest ways to introduce a self tuning of the processing parameters.

ACKONWLEDGEMENTS

We are indebted to Iris Voyat for running the program for geo-structural analysis and evaluating the results.

REFERENCES

Anderberg M. R. *Cluster analysis for applications*. Probability and Mathematical Statistics, Academic Press, NY, 1973.

Ferrero A.M., Forlani G., Voyat I.H., “Rock slope stability analysis based on photogrammetric surveys”, IX International Symposium on Landslides, Rio de Janeiro, 2004.

Fischler M., Bolles R., “Random sample consensus: a paradigm for model fitting with application to image analysis and automated cartography” *Commun. Assoc. Comp. Mach.*, Vol. 24:3, pp. 81-95, 1981.

Harrison, J. P., Reid, T. R.,“A semi-automated methodology for discontinuity trace detection in digital images of rock mass exposures.” *International Journal of Rock Mechanics & Mining Sciences* Vol. 37: pp 1073 – 1089, 2000.

Hartigan J. *Clustering Algorithms*, Wiley, 1975.

Roncella R., Forlani G., Remondino F., 2005. *Photogrammetry for geological applications: automatic retrieval of discontinuity orientation in rock slopes.* Proceedings of SPIE-IS&T Electronic Imaging, SPIE Vol. 5665 pp. 17-27.