

Unstructured point cloud semantic labeling using deep segmentation networks

A. Boulch, B. Le Saux and N. Audebert

ONERA - The French Aerospace Lab, FR-91761 Palaiseau, France

Abstract

In this work, we describe a new, general, and efficient method for unstructured point cloud labeling. As the question of efficiently using deep Convolutional Neural Networks (CNNs) on 3D data is still a pending issue, we propose a framework which applies CNNs on multiple 2D image views (or snapshots) of the point cloud. The approach consists in three core ideas. (i) We pick many suitable snapshots of the point cloud. We generate two types of images: a Red-Green-Blue (RGB) view and a depth composite view containing geometric features. (ii) We then perform a pixel-wise labeling of each pair of 2D snapshots using fully convolutional networks. Different architectures are tested to achieve a profitable fusion of our heterogeneous inputs. (iii) Finally, we perform fast back-projection of the label predictions in the 3D space using efficient buffering to label every 3D point. Experiments show that our method is suitable for various types of point clouds such as Lidar or photogrammetric data.

1. Introduction

The progress of 3D point cloud acquisition techniques and the democratization of acquisition devices have enabled the use of 3D models from real world in several economic fields such as building industry, urban planning or heritage conservation. Today's devices, like laser scanners or photogrammetry tools, allow the production of very large and precise point clouds, up to millions of points, structured or not. Meanwhile, the last years have seen the development of algorithms and methodologies in order to reduce the human intervention for two of the most common processing tasks with point clouds: first, surface reconstruction and abstraction, and second, object recognition and scene semantic understanding. However, these tasks are still a pending research topic and in applied fields, point cloud processing remains at least partly manual.

This work address the second issue: we aim at discovering the semantics of the scene, i.e. recognizing various classes of objects or content in the scene. In [BDM14], the semantic discovery of a scene is done using grammars on a 3D reconstructed model, so that the result is very dependent on the quality of the abstract model. Here, we adopt a different approach. Similarly to [HWS16, GKF09, LM12], we want to extract semantic information as soon as possible in the processing pipeline. As a matter of fact, knowing segmentation of the scene and the class of each object allows to direct the reconstruction according to each class: model or primitive fitting, regularity or symmetry constraints. More precisely, we aim at attributing a class label to each 3D point. In the image processing field, the similar task would be pixel wise labeling or semantic segmentation. Recent work on the subject focus on the design of ef-




Figure 1: Generation of 2D snapshots for semantic labeling in the image space by taking random camera positions in the 3D space.

ficient 3D descriptors by taking into account the neighborhoods of points [RHBB09, TSDS10]. We propose a different approach based on Convolutional Neural Networks (CNNs) and particularly on segmentation networks [LSD15, BKC15]. These networks reached the state of the art at image segmentation on different use cases such as KITTI [GLU12] or aerial images [ALSL16] on the ISPRS dataset [RSJ*12]. The originality of our approach is that our own

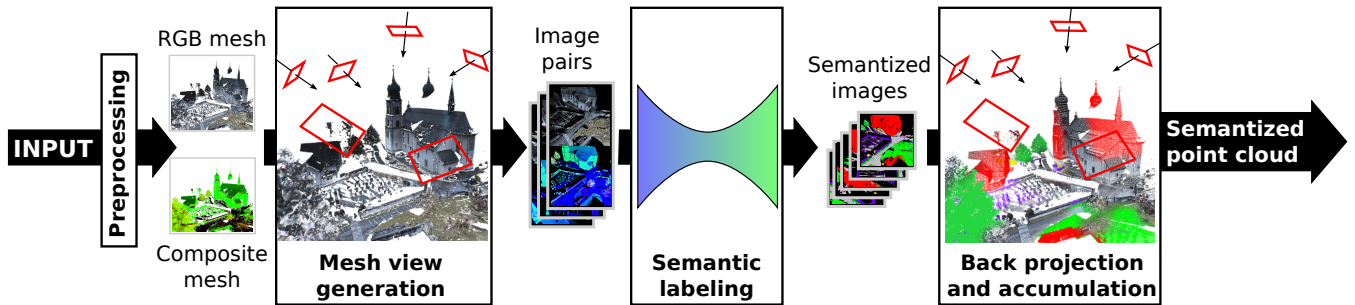


Figure 2: Approach work-flow.

features are simple 2D primitives: snapshots of the point-cloud. Then, we can do the labeling in a 2D image space (figure 1) where the segmentation networks proved to be very efficient. While the experiments presented in this paper are outdoor scenes, our labeling pipeline is generic and could be applied to various scenes and point cloud types.

Organization of the paper. The paper is organized as follows. The section 2 presents the related work on point cloud semantic labeling. The overview of our 4-step semantic labeling method can be found in section 3. Then the four next sections detail the main steps of the algorithm: section 4 explains the preprocessing of the 3D point-cloud required to take the snapshots according to the strategy exposed in section 5, the semantic labeling and data fusion pipeline based on convolutional networks is exposed in section 6 and point-cloud labeling is detailed in section 7. Finally, in section 8, we evaluate our segmentation method.

2. Related work

Semantic segmentation of point clouds is a well known problem in computational geometry and computer vision. Starting in the 90's, it gained in interest with the democratization of acquisition devices and reconstruction techniques [OK93]. The objective is to identify the class membership of each point. This problem is related to the 2D semantic segmentation, where the objective is to label each pixel of the image.

The early stages of semantic labeling for point cloud were mainly focused on aerial laser acquisition (Lidar). The objective was to discriminate building and roads from vegetation. A common approach is to discretize the point cloud on a regular grid to obtain a 2.5D elevation map authorizing to use image processing algorithms like in [HW97] where the authors use images filters or in [Maa99] for maximum likelihood classification. Other low level primitives, such as planes [BAvGT10], have also been used for bottom-up classification introduced in [HBA98] or [RB02].

In a more general context, low level shape extraction in point clouds has also been investigated. The Hough transform, originally designed for line extraction, was successfully adapted to 3D for plane extraction in [VGSR04]. [SWK07] proposes a generic RANSAC algorithm for geometric shape extraction in 3D point clouds. Hybrid shape extraction were investigated in [LKBH10, LM12] where the surfaces which fit geometric primitives are re-

placed with the corresponding abstract model while voids remain as triangular mesh.

Many algorithms for extraction of higher level semantic information were published in the recent years. In urban classification [HWS16, CGM09], classifying small objects like cars or street furniture [GKF09] and discriminating between roads and natural terrain become decisive at the smallest possible scale: point level [HWS16]. Most of the semantic labeling approaches rely on the same technique: designing the most discriminating features for the classification task. For example, in [CML04], the authors designed by hand a collection of expert features such as normalized height or luminance. Another approach is to create generic descriptor space to represent the points and their neighborhood in order to learn a supervised classifier. Among these descriptors, the spin images [JH99], the fast point feature histograms [RHBB09] or the signature histograms [TSDS10] may be the most popular. With respect to these approaches, we use much more simple features: 2D views of the point cloud.

By using a deep learning framework, it is possible to learn not only the classifier but also the feature representation. While deep neural networks are commonly used in image processing for classification or semantic labeling, there are only a few initiatives for semantic labeling in 3D [LBF14, WSK*15]. These approaches use a voxelization of the space to create 3D tensors in order to feed a convolutional neural network (CNN). However, using dense 3D representation for sparse input data consumes a lot of GPU memory and do not allows the use of large CNNs together with a refined voxel representation of the space. Even though there are great initiatives to efficiently reduce the memory cost on sparse data [Gra14], the direct 3D labeling is hardly tractable to personal computers and would require a whole server for training.

Apart from semantic segmentation, the application of deep learning in a 3D context knows an exponential interest, but the neural networks are mostly applied on 2D tensors. For example, in [LGK16], a deep framework is used to compute a metric for identifying architectural style distance between to building models. On a shape retrieval task, the authors of [SMKLM15] take several pictures of the 3D meshed object and then perform image classification using a deep network. Our approach has common features with this work: we generate snapshots of the 3D scene in order to use a 2D CNN with images as input. But unlike [SMKLM15] whose purpose is classification, i.e. giving a single label per 3D

shape, we compute dense labeling in the images and back project the result of the semantic segmentation to the original point cloud, which results in dense 3D point labeling.

3. Method overview

The core idea of our approach consists in transferring to 3D the very impressive results of 2D deep segmentation networks. It is based on the generation of 2D views of the 3D scene, as is someone was taking snapshots of the scene to sample it. The labeling pipeline is presented on figure 2. It is composed of four main processing steps: point-cloud preparation, snapshot generation, image semantic labeling and back projection of the segmentation to the original 3D space.

1. The **preprocessing** step aims at decimating the point cloud, computing point features (like normals or local noise) and generating a mesh.
2. **Snapshot generation**: from the mesh and the point attributes, we generate two types of views, Red-Green-Blue (RGB) and depth composite, by picking various camera positions (cf. Sec. 5).
3. **Semantic labeling** gives a label to each pair of corresponding pixels from the two input images. We use deep segmentation networks based on SegNet [BKC15] and fusion with residual correction [ALSL16].
4. Finally, we **project back to 3D** the semantized images. For each point of the mesh, we select its label by looking at the images where it is visible (cf. Sec. 7).


Point cloud properties In this work we assume our point clouds have a metric scale such that voxelization outputs have the same point density. We also consider as known the vertical direction to compute the normal deviation to this vector. As presented in section 8, it is also possible to use the pipeline without RGB information but performances are downgraded.

4. Point cloud preprocessing

The main issue for image generation when dealing with point clouds is the sparsity. When taking a snapshot, if the density of the point cloud is not sufficient one can see the points behind the observed structure. This leads to images which are difficult to understand, even for a trained human expert. To overcome this issue, we generate a basic mesh of the scene. The figure 3 shows the kind



Figure 3: Point cloud (left) and mesh (right) seen from the same point of view: dense representations help understanding the scene.



(a) RGB texture. (b) Depth composite texture.

Figure 4: Meshes for taking synthetic snapshots of the 3D scene.

of images we obtain with and without meshing. We now detail the algorithmic steps.

PointCloud decimation Point-clouds captured with ground lasers have varying point densities depending on the distance to the sensor. So, we first decimate the point cloud and get a lighter cloud so that subsequent processing can be applied in tractable times. To do that, we voxelize the scene, and keep the closest point to each voxel center (along with its class label at training time). In this paper, we chose a voxel size of 0.1m. It proved to produce relatively small point clouds while preserving most of the original features and shapes. Stronger decimations may lead to discarding small objects. In our experiments with semantic 3D, we reduce point cloud sizes from 20M/429M points to 0.4M/2.3M points.

Mesh generation The only a priori knowledge we have about our point-clouds is that they have homogeneous density due to decimation. For practical purposes, we chose the mesh generation algorithm from [MRB09] among many standard methods. Although it does not give any guaranty about the topology of the generated mesh, it is not a concern for our snapshot application. It requires as input a point-cloud with normals, which we estimate by using the available code from [BM16]. We now denote the mesh by $\mathcal{M} = (V, F)$ with V the set of vertices and F the faces.

Composite colors We aim at using both color and volume information for semantic labeling. To achieve that, we create two textures for the mesh (cf. Fig. 4). The more straightforward is the RGB texture, which takes the original point colors (cf. Fig. 4a). Then, we extract two generic features of point clouds: normal deviation to vertical and a noise estimation at a given scale. The normal deviation to the vertical at point p is:

$$\text{normdev}_p = \arccos(|\mathbf{n}_p \cdot \mathbf{v}|)$$

where \mathbf{n}_p is the normal vector and \mathbf{v} is the vertical vector. The noise at a given point p is an estimation of the spread of the points in its neighborhood.

$$\text{noise}_p = \frac{\lambda_2}{\lambda_0}$$

where λ_0 (resp. λ_2) is the highest singular value (resp. the lowest) obtained doing a principal component analysis estimation by singular value decomposition. Our *depth composite* texture encodes the normal deviation on the green channel and the local noise on

the red one. The blue channel remains empty at this point, but later will be filled with depth (i.e. distance to the camera).

5. View generation

Once the meshes are constructed, we want to produce the images for semantic labeling. We used an approach similar to [SMKLM15]. We load the model in a 3D mesh viewer and generate random camera positions and orientations to take various snapshots.

The camera parameters are generated according to two different strategies. First, in the *random* strategy, the camera center coordinates are randomly picked in the bounding box of the scene, with an altitude between 10 and 30 meters. The view direction is picked in a 45° cone oriented to the ground. To ensure the production of meaningful pictures, i.e. that the camera looks at the scene, we impose 20% of the pixels should correspond to actual points. Second, in the *multiscale* strategy, we pick a point of the scene, pick a line which goes through the point, and generate three camera positions on this line, oriented towards the point: thus ensuring each camera looks at the scene at various, increasing scale (allowing more and more details to be seen).

For each camera position, we generate three 224×224 -pixel images, as shown on figure 5. The first one is a snapshot of the RGB mesh (figure 5a) and reflects the real texture of the model. The second one is the depth composite image (figure 5b), made of surface normal orientation and noise completed with the depth to the camera. In order to do the back projection efficiently, we also generate an image where the color of each face of F is unique so that we know which face is visible (figure 5c). Finally for training or validation purposes, when ground truth is available, we create the corresponding label image (Fig. 5d).

6. Semantic Labeling

CNNs are feed-forward neural networks which make the explicit assumption that inputs are spatially organized. They are comprised of learnable convolution kernels stacked with non linear activations, e.g. ReLU ($\max(0, x)$). Those filters perform feature extraction in order to build an internal abstract representation of the input, optimized for later classification.

Several deep convolutional neural networks architectures exist for semantic labeling, usually derived from the Fully Convolutional Networks [LSD15]. Those models usually take RGB images in input and infer structured dense predictions by assigning a semantic class to every pixel of the image. In this paper, we use custom implementations of two network variants with a symmetrical encoder-decoder structure: SegNet and U-Net.

- SegNet [BKC15] is illustrated in figure 6a. The encoder part of SegNet is based on VGG-16 [SZ14], a deep CNN with 16 layers designed for image classification. Only the convolutional part is kept, while the fully connected layers are dropped. The decoder performs upsampling using the *unpooling* operation. During unpooling, the feature maps in the decoder are upsampled by placing the values into the positions given by the indices of the maximum during the symmetrical pooling in the encoder.

- U-Net [RFB15] is shown in figure 6b. Also based on VGG-16 for the encoder part, it uses a different trick for upsampling. It concatenates the feature maps of the decoder convolutional layers upsampled by duplication with the symmetrical feature maps in the encoder. Later convolutions blend both types of information.

As we extract both RGB and depth composite information from the dataset, we want to fuse the data sources to improve the accuracy of the model, compared to only one source. We use several fusion strategies in order to exploit the complementarity of the depth and RGB information. Therefore, two parallel 3-channels segmentation networks are trained, one on the RGB data, the other on the composite data. The experimented strategies are the following :

- Activation addition fusion, i.e. averaging of the two models (figure 6c). The predictions of the two SegNet are simply averaged pixel-wise.
- Prediction fusion using residual correction [ALSL16] (figure 6e). A very short (3 layers) residual network [HZRS15] is added at the end of the two SegNet. It takes in input the before last feature maps and learns a corrective term to apply to the averaged prediction.

Moreover, we also experiment early data fusion using a pre-processing CNN that projects the two data sources into a 3-channel common representation (figure 6d). We then use this projection as input of the traditional SegNet.

Compared to model averaging, using a neural network to learn how to fuse the two predictions should achieve better results, as it will be able to learn when to trust the individual sources based

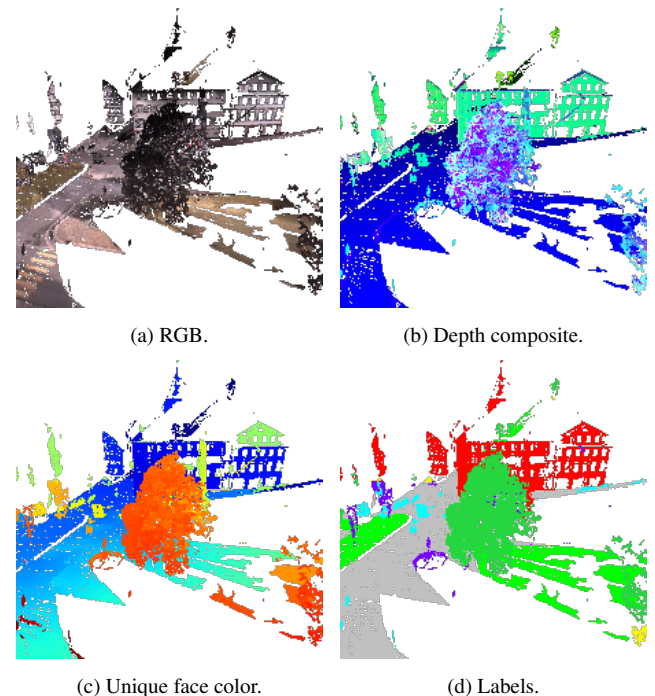


Figure 5: The various products of the preprocessing and view generation step.

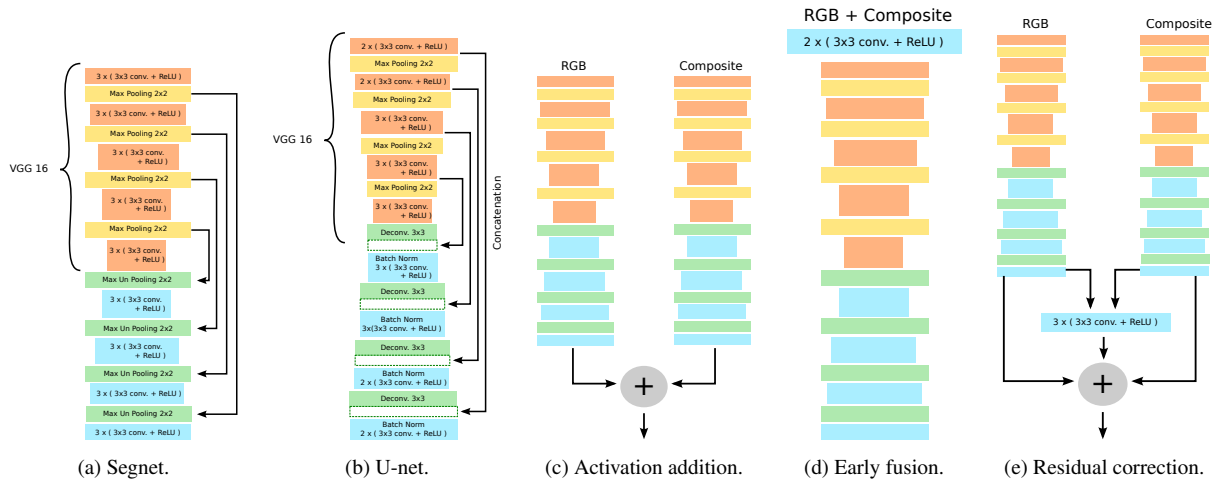


Figure 6: Various segmentation networks used in this paper: single-flow networks (a,b) vs. fusion networks (c,d,e)

on the context and the classes predicted. As an example, figure 7 presents a case of interest for fusion. The RGB prediction is wrong on the road. The network is fooled by the texture similar to the building one. On the other hand, the depth composite predicts the good label on the road but fails on the natural terrain where the steep slope has the geometric attributes of a building roof.

7. 3D back projection

This section presents how we project the pixel wise class scores obtained in section 6 on the original point cloud.

Projection to mesh. First we estimate the labels at each vertex of the mesh used to generate the images. Thanks to the unique-color-per-face images created at snapshot generation, we are able to quickly determine which faces are seen in each image pair and consequently the visible vertices of V . The score vector of the pixel is then added to the scores of each vertex of the face. This operation is iterated over all the images. Finally the vertex label is the class with the highest score.

Projection to the original point cloud. The second step is to project the labeled vertices to the original point cloud \mathcal{P} . We adopt a simple strategy. The label of a given point $p \in \mathcal{P}$ is the label of its nearest neighbor with label in V . For efficient computation, we build a k-d tree with V , and search for nearest neighbors. This allows not to load the whole \mathcal{P} , and avoid extensive memory al-




Figure 7: Mono input estimates: RGB (left) and composite (right).

location (particularly when dealing with hundreds of millions of points).

8. Experiments

In this section, we present the results of our experiments on semantic labeling of 3D point sets. We mainly experiment on the Semantic 3D dataset [HSL*16][semantic3d.net). The semantic-8 dataset is composed of 30 laser acquisitions (15 for training and 15 for testing) on 10 different scenes from various places and landscape types (rural, suburban, urban). The ground truth is available for the training set, and undisclosed for the test. There are 8 classes, namely: man-made terrain (gray), natural terrain (green), high vegetation (dark green), low vegetation (yellow), buildings (red), hard-scape (purple) scanning artefacts (cyan), cars (pink).

For quantitative evaluation, we use the same metrics as the dataset benchmark. It includes the overall accuracy (OA): $OA = \frac{T}{|\mathcal{P}|}$ where $|\mathcal{P}|$ is the size of the point cloud, and T is the number of true positive i.e the number of points that received the good label. We also use the intersection over union (IoU) per class: $IoU_c = \frac{T_c}{|\mathcal{P}_c \cup \hat{\mathcal{P}}_c|}$ where T_c is the number of points of class c correctly estimated, \mathcal{P}_c is the set of points with true label c and $\hat{\mathcal{P}}_c$ is the set of points with estimated class c . Finally the global average IoU (AIoU) is defined as: $IoU = \frac{1}{|C|} \sum_{c \in C} IoU_c$

8.1. Architecture and parameter choice

Dataset and training. In these experiments, we defined our own custom validation set by splitting the training set: 9 acquisitions for training and 6 for validation. For each training acquisition, we generated 400 image pairs, so that we optimize the deep networks with 3600 samples. We used a stochastic gradient descent with momentum (momentum is set to 0.9). The learning rate varies according to a step down policy starting at 0.01. It is multiplied by 0.2 every 30 epoch. The encoder part of SegNet is initialized with the VGG16 weights [SZ14].

At testing, we generate 500 views at 3 scales. For a point-cloud

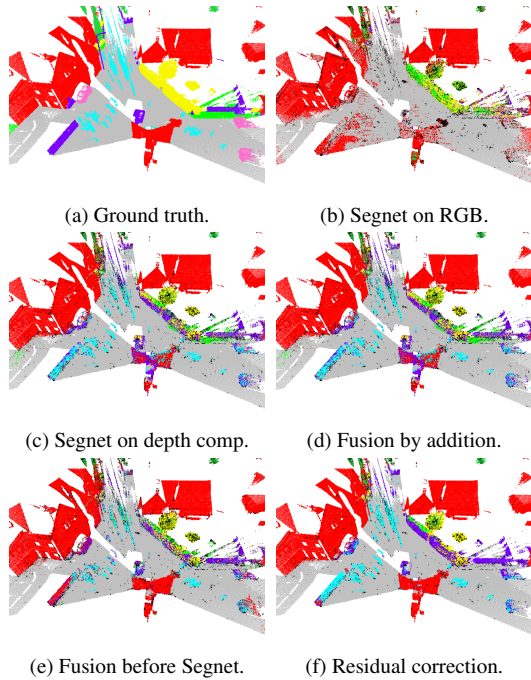


Figure 8: Same prediction view for the different fusion strategies.

of 30M points, the computation times are the following (with: CPU Xeon 3.5GHz, GPU TitanX Maxwell): pre-processing 25 min. (7 min. with normal estimation by regression); view-generation 7 min.; inference 1 min.; back-projection 8 min.; which sum up to 41 or 23 min. for the whole point-cloud semantization. Most sensitive parameters are the number of voxels (for point-cloud decimation) and the number of snapshots.

Fusion strategy choice. As explained in section 6, the different natures of the input images impose to define a fusion strategy. We quantitatively evaluate the different fusion options presented on figure 6. The results are presented on table 1a.

As a baseline, in the first result block, we trained two mono-input SegNets, taking as input the RGB or depth composite images. The composite network performs globally better except on buildings for which there is a great difference of texture compared to the rest of the scene. Moreover, depth composite images, that only contain geometric information, are not sensible to the texture of objects, so almost every vertical plane will be labeled as a building. This experiment shows that the two inputs are complementary and that the RGB network is not able to extrapolate the composite information only from the image texture.

The second result block of table 1a is dedicated to fusion strategies. Due to the high distribution difference in the prediction maps, composite-only totally overcomes the RGB prediction, i.e. the depth composite is most of the time confident while the RGB is more hesitating. As a result, the addition of prediction scores does not improve the results compared to depth composite only. A visual glimpse of the phenomenon is presented on figures 8c and 8d, the two images are almost similar.

Operating the fusion before labeling via SegNet should overcome this issue by melting the two signals at an early stage. As expected the results are visually improved (figure 8), particularly on natural terrain class, where the association of the texture and geometric features is discriminatory. However, the fusion step before SegNet is not optimal. VGG-16 takes a 3-channel image as input, and the two convolutions added before SegNet operates a dimension reduction that may cause information loss. Moreover, the different nature of the input makes it uncertain that information from both are compatible for fusion this early in the process. Finally the best results are obtained by the residual correction network. The compromise between the fusion after Segnet (addition) and a more refined fusion using convolution (previous case) is successful. The residual correction compensates the difference of the two outputs, resulting in an increase of the performances on almost all classes.

For comparison with existing methods we confront our approach to full semantic-8 dataset. We present the results in table 1b. The three other methods are the publicly available results. [MZWLS14] is method for aerial images based segmentation on images descriptors and an energy minimization on a conditional random field. In [HWS16], the authors use a random forest classifier trained on multi-scale 3D features taking into account both surface and context properties. Harris Net is not described on the result board, but from its name, we assume a method based on 3D Harris point extraction followed by a classification using a deep framework. We present the results of two methods, SegNet with a purely random set of images, and a U-Net with zoom on snapshot strategy. To our knowledge the two networks performs equally and the main difference reside in the snapshot strategy. At redaction time, our U-Net took the first place in the leaderboard for global scores, average IoU and overall accuracy. Looking at the per class IoU, we take the lead on six out of eight categories. Among them, the performances on natural terrain, scanning artifacts and cars are drastically increased. On man-made terrain and buildings, we place second with a comparable score as [HWS16] and Harris Net. The use of the zoom strategy greatly improves the score on cars and scanning artifacts. The reason is that compared to the random strategy, the training dataset (and the test dataset) contains more images with small details, which makes them possible to segment. The only relative failure of the deep segmentation networks are the scanning artifacts and the hardscape classes. Even though we place first on these categories, the IoU score is low: we discuss this in section 8.3.

8.2. Photogrammetric point clouds

In order to evaluate the capacity of our method to be transferable, we also experiment on photogrammetric data. The figure 9 presents a reconstruction of Mirabello's church destroyed after an earthquake in 2012 in Italy. We followed the same process as for the laser data. The network used for semantic labeling is the one trained on the full semantic-8 training set. The results are fairly encouraging. Most of the visual error concentrates on ground classes and high vegetation. A lot of ground is covered by rubble coming from the destroyed building. Due to the chaotic structure of the debris, it is recognized as natural terrain. Part of the rooftops are also wrongly labeled the same way. We interpret this as a consequence of the fact that our training set contains only ground laser acquisitions so only

Method	AIoU	OA	IoU1	IoU2	IoU3	IoU4	IoU5	IoU6	IoU7	IoU8
SegNet RGB	0.28	0.749	0.853	0.097	0.483	0.075	0.69	0.042	0.0	0.0
SegNet Depth Comp.	0.326	0.763	0.902	0.342	0.597	0.013	0.503	0.178	0.066	0.003
SegNet add.	0.312	0.762	0.895	0.237	0.573	0.029	0.522	0.172	0.067	0.003
SegNet before	0.336	0.763	0.898	0.569	0.452	0.021	0.510	0.179	0.051	0.009
SegNet Res.	0.427	0.805	0.948	0.739	0.763	0.024	0.710	0.133	0.097	0.0

(a) Comparison of deep segmentation networks on Semantic 3D, custom test set.

Method	AIoU	OA	IoU1	IoU2	IoU3	IoU4	IoU5	IoU6	IoU7	IoU8
Graphical models [MZWLS14]	0.391	0.745	0.804	0.661	0.423	0.412	0.647	0.124	0.000	0.058
Random forest [HWS16]	0.494	0.850	0.911	0.695	0.328	0.216	0.876	0.259	0.113	0.553
Harris Net	0.623	0.881	0.818	0.737	0.742	0.625	0.927	0.283	0.178	0.671
Ours SegNet Rand. Snap.	0.516	0.884	0.894	0.811	0.590	0.441	0.853	0.303	0.190	0.050
Ours U-Net Multiscale Snap.	0.674	0.910	0.896	0.795	0.748	0.561	0.909	0.365	0.343	0.772

(b) Semantic 3D results on full test set.

IoU: intersection over union (per class), AIoU: average intersection over union, OA: overall accuracy. Classes 1: man-made terrain, 2: natural terrain, 3: high vegetation, 4: low vegetation, 5: buildings, 6: hardscape, 7: scanning artefacts, 8: cars.

Table 1: Quantitative results on Semantic 3D.

sloping roofs are present at training time. When confronted to roofs with small inclination, the network is misled to a ground class. Finally, high vegetation labels appear on destroyed parts which are still standing. This is mainly due to the high noise estimation (red channel of the depth composite image) which is incompatible with building.

8.3. Limitations and perspectives

Even though the proposed approach obtains the best performances on the semantic-8 leader board there are still issues to overcome. First, a non-exhaustive training set influences the results: for example missing architectural elements or samples may explain the relatively low scores on hardscape and scanning artifacts. For a more generic pipeline, one should use a more diversified training set. A second field of future investigation is postprocessing the results to remove the outliers by regularization. For example, we could enforce the volumetric consistency of labels in a neighborhood or impose constraints on points belonging to a common extracted shape. Another question is the suitability of the method for point-clouds obtained by accumulating data of low-cost range cameras. They can be dense enough, but will be noisier than laser point-clouds. Finally, a promising line of investigation is to perform data augmentation by using data from other sources. For example, synthetically generated images could be added in the training set, or scenes could be augmented with 3D models of small objects like cars. In addition to modifying the proportion of given classes, it increases the variability of the scenes (more configurations) and consequently avoids overfitting which leads to a more generic framework.

9. Conclusion

We have presented a new and efficient framework for semantic labeling of 3D point clouds using deep segmentation neural networks. We first generate RGB and geometric composite images of the scene. These pairs are the inputs of our network architectures for semantic segmentation. Several strategies for data fusion were

investigated, and among them segmentation network with residual correction proved to perform the best. Finally, image segmentation were aggregated on the 3D model to give to each point a label. We experimented on both laser scans and photogrammetric reconstruction. The method was evaluated against the semantic-8 dataset and obtained the best performances in the leader board concerning the global measurements and several individual classes. We also got encouraging results of transferring networks trained on laser acquisition to photogrammetric data. Although we obtain good performances, several fields of investigation remain such as data augmentation or images generation strategies to improve the scores on small and rare classes.

Implementation details

The manipulation of point clouds, i.e the preprocessing, the image creation and the back projection was implemented using Python and C++; with PCL and the 3D viewer of pyqtgraph.org. The neural networks were implemented using Tensorflow.

Acknowledgments

The research of A. Boulch and B. Le Saux has received funding from the European Union's Seventh Framework Programme (FP7/2007-2013) under grant agreement no. FP7-SEC-607522 (Inachus Project). N. Audebert's work is funded by ONERA-TOTAL research project Naomi.

References

- [ALSL16] AUDEBERT N., LE SAUX B., LEFÈVRE S.: Semantic Segmentation of Earth Observation Data Using Multimodal and Multi-scale Deep Networks. In *ACCV* (Taipei, Taiwan, Nov. 2016). 1, 3, 4
- [BAVGT10] BUGHIN E., ALMANSA A., VON GIOI R. G., TENDERO Y.: Fast plane detection in disparity maps. In *ICIP* (2010), IEEE, pp. 2961–2964. 2




Figure 9: Semantic labeling of photogrammetric data.

- [BDM14] BOULCH A., DE LA GORCE M., MARLET R.: Piecewise-planar 3D reconstruction with edge and corner regularization. In *Computer Graphics Forum* (2014), vol. 33, pp. 55–64. 1
- [BKC15] BADRINARAYANAN V., KENDALL A., CIPOLLA R.: SegNet: A Deep Convolutional Encoder-Decoder Architecture for Image Segmentation. *arXiv preprint arXiv:1511.00561* (2015). 1, 3, 4
- [BM16] BOULCH A., MARLET R.: Deep Learning for Robust Normal Estimation in Unstructured Point Clouds. *Computer Graphics Forum* (2016). 3
- [CGM09] CHEHATA N., GUO L., MALLET C.: Airborne Lidar feature selection for urban classification using random forests. *Int. Archives Photogramm. Remote Sens. Spat. Inf. Sci* 38, Part 3 (2009), W8. 2
- [CML04] CHARANIYA A. P., MANDUCHI R., LODHA S. K.: Supervised parametric classification of aerial Lidar data. In *CVPRW* (2004), IEEE, pp. 30–30. 2
- [GKF09] GOLOVINSKIY A., KIM V. G., FUNKHOUSER T.: Shape-based recognition of 3D point clouds in urban environments. *ICCV* (Sept. 2009). 1, 2
- [GLU12] GEIGER A., LENZ P., URTASUN R.: Are we ready for autonomous driving? the KITTI vision benchmark suite. In *CVPR* (2012), IEEE, pp. 3354–3361. 1
- [Gra14] GRAHAM B.: Spatially-sparse convolutional neural networks. *CoRR abs/1409.6070* (2014). 2
- [HBA98] HAALA N., BRENNER C., ANDERS K.-H.: 3D urban GIS from laser altimeter and 2D map data. *International Archives Photogramm. Remote Sens.* 32 (1998), 339–346. 2
- [HSL*16] HACKEL T., SAVINOV N., LADICKY L., WEGNER J.-D., SCHINDLER K., POLLEFEYS M.: Large-scale point cloud classification benchmark. In *CVPR/ Large Scale 3D Data Workshop* (2016). 5
- [HW97] HUG C., WEHR A.: Detecting and identifying topographic objects in imaging laser altimeter data. *International archives Photogramm. Remote Sens.* 32, 3 SECT 4W2 (1997), 19–26. 2
- [HWS16] HACKEL T., WEGNER J. D., SCHINDLER K.: Fast semantic segmentation of 3D point clouds with strongly varying density. *ISPRS Ann. Photogramm. Remote Sens. Spat. Inf. Sci* 3 (2016), 177–184. 1, 2, 6, 7
- [HZRS15] HE K., ZHANG X., REN S., SUN J.: Deep residual learning for image recognition. *arXiv preprint arXiv:1512.03385* (2015). 4
- [JH99] JOHNSON A. E., HEBERT M.: Using spin images for efficient object recognition in cluttered 3D scenes. *IEEE PAMI* 21, 5 (1999), 433–449. 2
- [LBF14] LAI K., BO L., FOX D.: Unsupervised feature learning for 3D scene labeling. In *ICRA* (2014), IEEE, pp. 3050–3057. 2
- [LGK16] LIM I., GEHRE A., KOBELT L.: Identifying style of 3D shapes using deep metric learning. *Computer Graphics Forum* 35, 5 (2016), 207–215. 2
- [LKBH10] LAFARGE F., KERIVEN R., BRÉDIF M., HIEP V. H.: Hybrid multi-view reconstruction by jump-diffusion. In *CVPR* (2010), IEEE, pp. 350–357. 2
- [LM12] LAFARGE F., MALLET C.: Creating large-scale city models from 3D-point clouds: a robust approach with hybrid representation. *Int. journal of computer vision* 99, 1 (2012), 69–85. 1, 2
- [LSD15] LONG J., SHELHAMER E., DARRELL T.: Fully Convolutional Networks for Semantic Segmentation. In *CVPR* (2015), pp. 3431–3440. 1, 4
- [Maa99] MAAS H.-G.: The potential of height texture measures for the segmentation of airborne laser scanner data. In *21st Canadian symp. on remote sensing* (1999), pp. 154–161. 2
- [MRB09] MARTON Z. C., RUSU R. B., BEETZ M.: On Fast Surface Reconstruction Methods for Large and Noisy Datasets. In *ICRA* (Kobe, Japan, May 12-17 2009). 3
- [MZWLS14] MONTOYA-ZEGARRA J. A., WEGNER J. D., LADICKY L., SCHINDLER K.: Mind the gap: modeling local and global context in (road) networks. In *GCP* (2014), Springer, pp. 212–223. 6, 7
- [OK93] OKUTOMI M., KANADE T.: A Multiple-Baseline Stereo System. *IEEE PAMI* 15(4) (1993), 353–363. 2
- [RB02] ROTTENSTEINER F., BRIESE C.: A new method for building extraction in urban areas from high-resolution Lidar data. *Int. Archives Photogramm. Remote Sens. Spat. Inf. Sci* 34, 3/A (2002), 295–301. 2
- [RFB15] RONNEBERGER O., FISCHER P., BROX T.: U-Net: Convolutional networks for biomedical image segmentation. In *MICCAI* (Munich, 2015), pp. 234–241. 4
- [RHBB09] RUSU R. B., HOLZBACH A., BLODOW N., BEETZ M.: Fast geometric point labeling using conditional random fields. In *IROS* (2009), IEEE, pp. 7–12. 1, 2
- [RSJ*12] ROTTENSTEINER F., SOHN G., JUNG J., GERKE M., BAILLARD C., BENITEZ S., BREITKOPF U.: The ISPRS benchmark on urban object classification and 3D building reconstruction. *ISPRS Ann. Photogramm. Remote Sens. Spat. Inf. Sci* 1 (2012), 293–298. 1
- [SMKLM15] SU H., MAJI S., KALOGERAKIS E., LEARNED-MILLER E.: Multi-view convolutional neural networks for 3D shape recognition. In *ICCV* (2015), pp. 945–953. 2, 4
- [SWK07] SCHNABEL R., WAHL R., KLEIN R.: Efficient RANSAC for point-cloud shape detection. In *Computer graphics forum* (2007), vol. 26, Wiley Online Library, pp. 214–226. 2
- [SZ14] SIMONYAN K., ZISSERMAN A.: Very deep convolutional networks for large-scale image recognition. *CoRR abs/1409.1556* (2014). 4, 5
- [TSDS10] TOMBARI F., SALTI S., DI STEFANO L.: Unique signatures of histograms for local surface description. In *ECCV* (Hersonissos, Crete, 2010), Springer, pp. 356–369. 1, 2
- [VGSR04] VOSSELMAN G., GORTE B. G., SITHOLE G., RABBANI T.: Recognising structure in laser scanner point clouds. *Int. Archives Photogramm. Remote Sens. Spat. Inf. Sci* 46, 8 (2004), 33–38. 2
- [WSK*15] WU Z., SONG S., KHOSLA A., YU F., ZHANG L., TANG X., XIAO J.: 3D shapenets: A deep representation for volumetric shapes. In *CVPR* (Boston, USA, 2015), pp. 1912–1920. 2