Machine-learning classification of aerial photogrammetric 3D point clouds

Carlos BECKER, Elena ROSINSKAYA, Nicolai HANI, Emmanuel D'ANGELO, Christoph STRECHA, Switzerland

Key words: Semantic Classification, LiDAR, Point Clouds, Photogrammetry, Digital Terrain Model.

SUMMARY

Pix4D presents a powerful method to extract per-point semantic class labels from photogrammetry data. Labelling this kind of data is important for tasks such as environmental modelling, object classification, and scene understanding. Unlike previous point cloud classification methods that rely exclusively on geometric features, we show that merging the color information from each pixel value with geometry features yields a significant increase in accuracy in detecting semantic classes.

Machine-learning classification of aerial photogrammetric 3D point clouds (9340)
Carlos Becker (Switzerland), Nicolai Häni (USA), Elena Rosinskaya, Emmanuel d'Angelo and Christoph Strecha (Switzerland)

Machine-learning classification of aerial photogrammetric 3D point clouds

Carlos Becker, Elena Rosinskaya, Nicolai Häni, Emmanuel d'Angelo, Christoph Strecha, Switzerland

1. INTRODUCTION

Extraction of semantic information from point clouds enables us to understand a scene, classify objects and generate high-level models with CAD-like geometries from them. It can also provide a significant improvement to existing algorithms, such as those used to construct Digital Terrain Models (DTMs) from Digital Surface Models (DSMs) (Unger et al., 2009). With the growing popularity of laser scanners, the availability of drones as surveying tools, and the rise of commercial photogrammetry software capable of generating millions of points from images, the need for fully automated extraction of semantic information from this kind of data increases as well. Although some of the commercial photogrammetry software available today offer tools such as automated DTM extraction (Pix4Dmapper, 2017, Photoscan, 2017), semantic classification is typically left to specialized software packages (eCognition, 2017, GlobalMapper, 2017) that rely on 2.5D orthomosaics and DSMs as an input.

The need for semantic modeling of 3D point data has inspired many research and application engineers to model specific structures. Often the proposed solutions were handcrafted to the application at hand: buildings have been modeled by using common image processing techniques such as edge detection (Haala et al., 1998, Brenner, 2000) or by fitting planes to point clouds (Rusu et al., 2007) or even building models (Descombe, 2007); road networks have been modeled by handcrafted features and DTM algorithms used heuristics about the size of objects. While successful and valuable, these approaches are inherently limited since they cannot be easily applied to detect new classes of objects. The huge boost in the performance of machine learning algorithms over the last years allows for more flexible and general learning and classification algorithms. Therefore we focus here on machine learning techniques that will allow the users to detect objects categories of their own choice.

In this paper we present a method to classify aerial photogrammetry point clouds. Our approach exploits both geometric and color information to assign to individual points class labels from the LAS standard: buildings, terrain, high vegetation, roads, human made objects (HMO) or cars. We show that incorporating color information yields a significant increase in accuracy compared to previous point cloud classification methods that rely exclusively on geometric features.
We evaluate our approach on four challenging datasets and show that off-the-shelf machine learning techniques together with our new features result in highly accurate and efficient classifiers that generalize well to unseen data. The datasets used for evaluation are publicly available at https://pix4d.com/research. Moreover, we show that our classification approach can be used to generate accurate Digital Terrain Models, without the need for hand-designed

heuristics such as Maximally Stable Extremal Regions (MSER) detection on a Digital Surface Model.

2. RELATED WORK

Methods used to extract semantic information from point clouds can be split into two groups: those that try to segment coherent objects from a scene, and those that focus on assigning an individual class label to each point. Early works using the first approach often converted the point data into a regular 2.5D height grid so that standard image processing techniques, e.g., edge detection, can be applied (Haala et al., 1998, Haala and Brenner, 1999, Wang and Schenk, 2000). A scan line based approach (Sithole and Vosselman, 2003) was proposed for structure detection in urban scenes. Building extraction approaches typically use geometric primitives during the segmentation step. A multitude of such primitives has been proposed, both in 2D, such as planes and polyhedral (Vosselman et al., 2001, Dorninger and Nothegger, 2007), and in 3D (Lafarge and Mallet, 2012, Xiao and Furukawa, 2014). In (Rusu et al., 2007) the authors fit sampled parametric models to the data for object recognition. Similarly, (Oesau et al., 2016) investigates supervised machine learning techniques to represent small indoor datasets with planar models for object recognition.

The second type of methods assign a label to each point in the point cloud. Typically this is done with supervised machine learning techniques. Binary classification has been explored in environmental monitoring to extract road surfaces (Shu et al., 2016), tree species (Böhm et al., 2016, Liu and Böhm, 2015), land cover (Zhou et al., 2016), and construction sites (Xu et al., 2016). Several other authors employed a multiclass setting to classify multiple types of objects and structures (Brodu and Lague, 2012, Weinmann et al., 2015a, Hackel et al., 2016), which we adopt in this paper. In particular, we follow the work of (Weinmann et al., 2013), which introduced local geometric features that were used to train a Random Forest (RF) classifier for single terrestrial LiDAR scans. Their set of features was extended later by (Hackel et al., 2016). Examples of other feature sets used in the point classification context are Fast Point Feature Histogram (FPFH) (Rusu et al., 2009) or Color Signature of Histogram of Orientations (SHOT) (Tombari et al., 2010). All these methods use handcrafted features that can be considered suboptimal when compared to more recent deep learning techniques (Hu and Yuan, 2016, Qi et al., 2016), which learn features directly on image or point cloud data. Those approaches have not been considered here, since they require large computational power and large batches of data to train the classifier, and may be restrictive at prediction time, depending on the hardware available.

The ambiguity of the classification task can be minimized by modeling also the spatial correlations between the different class labels. Spatial priors are used in (Shapovalov and Velizhev, 2011) to classify LiDAR data and in (Niemeyer et al., 2014) the authors apply Conditional Random Field (CRF) priors to model different probabilities that neighboring labels can have. While those methods show reasonable classification improvements, they are computationally expensive and not easy to parallelize.

In this paper we extend the work on geometric features by (Weinmann et al., 2013, Hackel et al., 2016) and show that incorporating color information provides a significant boost in

Carlos Becker (Switzerland), Nicolai Häni (USA), Elena Kosinskaya, Emmanuel d'Angelo and Christoph Strecha
(Switzerland)

prediction accuracy, while keeping a low computational load at prediction time. In the following sections we describe our method and present the results obtained on four photogrammetry datasets.

3. METHOD

Our approach combines geometric and color features that are fed to a classifier to predict the class of each point in the point cloud. The geometric features are those introduced in (Hackel et al., 2016), which are computed at multiple scales to capture details at varying spatial resolutions, as explained shortly further. To exploit color information, we compute additional color features, based on the color of the respective point and its neighbors.

3.1 Geometric Features

To compute the features for a single scale for a point x, we first obtain its neighboring points $S_x = \{p_1, p_2, \dots, p_k\}$. This set is used to compute a local 3D structure covariance tensor

$$C_{\mathbf{x}} = \frac{1}{k} \sum_{i=1}^{k} (\mathbf{p}_i - \hat{\mathbf{p}})(\mathbf{p}_i - \hat{\mathbf{p}})^T$$

where

$$\hat{\mathbf{p}} = \arg\min_{\mathbf{p}} \sum_{i=1}^{k} ||\mathbf{p}_i - \mathbf{p}||$$

is the medoid of Sx.The eigenvalues $\lambda_1 \geq \lambda_2 \geq \lambda_3 \geq 0$ are unit-sum normalized and the corresponding eigenvectors $\mathbf{e_1}, \mathbf{e_2}, \mathbf{e_3}$ of C_x are used to compute the local geometry features shown in Table 1.

Sx.The eigenvalues $\lambda_1 \geq \lambda_2 \geq \lambda_3 \geq 0$ are unit-sum normalized and the corresponding eigenvectors $\mathbf{e_1}, \mathbf{e_2}, \mathbf{e_3}$ of C_x are used to compute the local geometry features shown in Table 1.

We have slightly changed the geometry feature set from (Hackel et al., 2016) and removed the sum of eigenvalues because it is constant since the eigenvalues are normalized to unit sum. In addition there are the first and second order moments of the point neighborhood around the eigenvectors which help to identify edges and occlusions. Besides the features based on the eigenvalues and eigenvectors of C_x, features based on the z coordinate of the point are used to increase their discriminative power.

3.2 Multi-scale Pyramid

To incorporate information at different scales we follow the multi-scale approach of (Hackel et al., 2016), which has shown to be more computationally efficient than that of (Weinmann et

Carlos Becker (Switzerland), Nicolai Häni (USA), Elena Rosinskaya, Emmanuel d'Angelo and Christoph Strecha
(Switzerland)

Covariance	Omnivariance	$(\lambda_1 \cdot \lambda_2 \cdot \lambda_3)^{\frac{1}{3}}$
	Eigenentropy	$-\sum_{i=1}^{3} \lambda_i \cdot \ln(\lambda_i)$
	Anisotropy	$(\lambda_1 - \lambda_3)/\lambda_1$
	Planarity	$(\lambda_2 - \lambda_3)/\lambda_1$
	Linearity	$(\lambda_1 - \lambda_2)/\lambda_1$
	Surface variation	λ_3
	Scatter	λ_3/λ_1
	Verticality	$1 - \lvert\langle [0,0,1], \mathbf{e}_3 \rangle\rvert$
Moments	1^{st} order, 1^{st} axis	$\sum_{\mathbf{p} \in \mathcal{S}_\mathbf{x}} \langle \mathbf{p} - \hat{\mathbf{p}}, \mathbf{e}_1 \rangle$
	1^{st} order, 2^{st} axis	$\sum_{\mathbf{p} \in \mathcal{S}_\mathbf{x}} \langle \mathbf{p} - \hat{\mathbf{p}}, \mathbf{e}_2 \rangle$
	2^{st} order, 1^{st} axis	$\sum_{\mathbf{p} \in \mathcal{S}_\mathbf{x}} \langle \mathbf{p} - \hat{\mathbf{p}}, \mathbf{e}_1 \rangle^2$
	2^{st} order, 2^{st} axis	$\sum_{\mathbf{p} \in \mathcal{S}_\mathbf{x}} \langle \mathbf{p} - \hat{\mathbf{p}}, \mathbf{e}_2 \rangle^2$
Height	Vertical Range	$z_{\max\{\mathcal{S}_\mathbf{x}\}} - z_{\min\{\mathcal{S}_\mathbf{x}\}}$
	Height below	$z_\mathbf{p} - z_{\min\{\mathcal{S}_\mathbf{x}\}}$
	Height above	$z_{\max\{\mathcal{S}_\mathbf{x}\}} - z_\mathbf{p}$
Color	Point color	$[H_\mathbf{x}, S_\mathbf{x}, V_\mathbf{x}]$
	Neighborhood colors	$\frac{1}{\lvert\mathcal{N}_\mathbf{x}(r)\rvert} \sum_{\mathbf{p} \in \mathcal{N}_\mathbf{x}(r)} [H,S,V]_\mathbf{p}$

Table 1: Our set of geometric (top) and color features (bottom) computed for points in local neighborhood Sx . Geometric features are based on eigenvalues of the local structure tensor, moments around the corresponding eigenvectors, height differences in Sx . Color features include HSV space color of the point of interest and its neighborhood.

3.3 Color Features

To increase the discriminative power of the feature set we combine the geometric features introduced above with color features. Our color features are computed in the HSV color space first introduced by (Smith, 1978), since the analysis of the Pearson product-moment correlation coefficient and the Fisher information of our training data showed that we should expect higher information gain from the HSV over RGB color space.

Carlos Beckert (Switzerland), Nicolai Hani (USA), Elena Kosmatopoulou, Emmanuel d'Angelo and Christoph Strecha
Machine Learning Classification for a Point Cloud Photogrammetric 3D Point Cloud (8140)

Besides the HSV color values of the point itself, we compute the average color values of the points within a certain radius around the query point in the original point cloud (i.e. not downsampled). We experimented with radii of 0.4 m, 0.6 m and 0.9 m to balance between classification speed and accuracy.

3.4 Training and classification

We use *Gradient Boosted Trees* (GBT) (Friedman et al., 2001) with a leave-one-out evaluation methodology: we train on three point clouds and test on the remaining one.. GBT can generate conditional probabilities and are applicable to multi-class classification problems. They are easily parallelized and are available as reusable software packages in different programming languages. To speed up prediction we implement a basic early stopping scheme on top of GBT: at prediction time and for each sample, we compute the margin of the most-voted class every e_n evaluated trees. If the margin is larger than a given threshold we assume that the classifier is confident enough about this sample, and therefore there is no need to evaluate the remaining trees in the ensemble. We will show later that this scheme speeds up prediction, improving user experience in interactive applications.

3.5 Implementation Details

We implemented our software in C++ to ease its later integration into the Pix4DMapper software. For prototyping and evaluation we used Julia (Bezanson et al., 2014). For fast neighbor search we used the header-only nanoflann library (https://github.com/jlblancoc/nanoflann) which implements a kd-tree search structure. We parallelized training and prediction, reducing computation times significantly. For GBT we used Microsoft's LightGBM (https://github.com/Microsoft/LightGBM).

4. EVALUATION

4.1 Datasets.

Table 2 shows the characteristics of the datasets employed for evaluation. The Paris-rue-Madame dataset (Serna et al., 2014) does not contain color information and was solely used to verify that our geometric features perform as well as those of (Hackel et al., 2016). Our main interest is the aerial photogrammetry and the four last datasets of Table 2. The images were processed with Pix4Dmapper to obtain their respective dense point clouds that were used as the input for our approach. Note that the GSD varies significantly between datasets. A 3D visualization is presented in Fig. 3(a). Moreover, each dataset contains different types of objects and terrain surfaces as shown in Table 3. For example, while all datasets contain roads, cropland only appears in one of them. This table will be useful later to analyze the performance of our approach on each dataset. We have made three photogrammetry datasets publicly available at https://pix4d.com/research.

dataset. Finally, although high vegetation appears in the top-3 misclassified classes in Table 5, this is mostly due to ambiguities in the ground truth: some small bushes were manually labeled as ground, while the classifier predicts them as high vegetation.

The classifier performs very well on the Building dataset. The highest error is predicting buildings as high vegetation or human-made objects. This dataset has the lowest GSD (or highest resolution), and facades of the buildings are reconstructed well. This is not the case for the other three datasets with higher GSD, where only few facade points are available. We hypothesize that the classifier is confused with the facades, finding the vegetation or human-made object to be the closest match.

On the Cadastre dataset the classifier predicts vast amounts of ground and vegetation points as buildings and human-made objects, leading to a very high error rate. This result is expected considering Table 3, as the Cadastre dataset contains hills and non-flat ground surfaces, which are not present in any of the other two datasets. Thus, the classifier confuses points in the regions of inclined ground with other classes that are closer in feature space to the training data (e.g. building roofs present a slope that resembles the properties of the points on a hill).

Our approach performs very well on the Rural dataset, obtaining the lowest misclassification

(a) Original data.

(b) Classification with geometry features only.

(c) Classification with geometry + color features.

Figure 1: Qualitative results obtained with our approach on the Ankeny dataset, using the other three datasets fortraining. We used neighbor color features within a 0.6-meter radius neighborhood and the Gradient Boosted Treesclassifier. Incorporating color information into the classifier results improves classification, particularly for the roadsbetween buildings.

Carlos Becker (Switzerland), Nicolai Häni (USA), Elena Kosinskaya, Emmanuel d'Angelo and Christoph Strecha
(Switzerland)

error among all datasets of 6%. Most of this error is due to vegetation being classified as ground, which is partly because of ambiguities in the ground truth: it is hard to disambiguate high vegetation from ground near the border of a forest.

The analyses above highlight the importance of reliable and varied training data, in that it should resemble the unseen data on which the classifier will be applied, e.g different landscapes, seasons, shapes of buildings, etc.

Dataset		True Label	Predicted Label						Overall error
Training	Testing		Ground	High vegetation	Building	Road	Car	HMO	
Building		Ground	56%	2.3%	0.3%	**10%**	0.2%	0.6%	12.1%
Cadastre	**Ankeny**	High vegetation	**2.7%**	6.7%	0.1%	0.1%	0.1%	0.1%	3.0%
Rural		HMO	**0.3%**	0.1%	0.1%	0.1%	0.2%	0.3%	0.9%
Ankeny		Road	**4.2%**	0.1%	0.1%	32%	0.1%	0.4%	4.9%
Cadastre	**Building**	Building	0.2%	**1.8%**	30%	0.4%	0.1%	1.7%	4.2%
Rural		High vegetation	**2.0%**	10%	0.2%	0%	0%	0.2%	2.5%
Ankeny		Ground	36%	4.1%	4.1%	**5.1%**	1.4%	1.3%	14.8%
Building	**Cadastre**	Road	**3.0%**	0.7%	0.4%	15.8%	0.2%	0.9%	5.1%
Rural		High vegetation	0.6%	9.1%	**2.1%**	0.1%	0.0%	0.8%	3.7%
Ankeny		High vegetation	**2.9%**	50%	0.1%	0.02%	0%	1.2%	3.2%
Building	**Rural**	Ground	42%	**2.4%**	0.02%	0.2%	0%	0%	2.7%
Cadastre		Road	**0.2%**	0%	0%	1.3%	0%	0%	0.2%

Table 5: Confusion matrix for the top-3 misclassified classes. In bold we highlighted the misclassification error corresponding to the class with which the true label is confused the most. Results obtained for the G + $C_{N\,(0.6)}$ features, training on two datasets and testing on the remaining third one. Percentages are with respect to the total number of points in the

4.5 Qualitative Results

Figure 1 shows a 3D view of the Ankeny dataset and the respective classified point cloud obtained when using geometric features only, as well as with our approach. Overall the results are very satisfying, especially when one considers the heterogeneity of the different datasets, as discussed earlier.

4.5.1 Timings

In Table 6 can be seen the speed up obtained with early stopping. We set $e_n = 20$ and vary the early stopping threshold e_t. A threshold of $e_t = 1.5$ yields a speed up of 2 to 3.5 times compared to not using early stopping, with a very small error increment of 1%.

Phase	Ankeny	Building	Cadastre	Rural
Geom. feature computation	41	14	28	85

Carlos Becker (Switzerland), Nicolai Häni (USA), Elena Rosinskaya, Emmanuel d'Angelo and Christoph Strecha
(Switzerland)

heuristics. The classification method presented in this paper is a part of Pix4Dmapper Pro. Earlier we mentioned that access to properly labeled training data that represents aerial photogrammetry point clouds is limited. To overcome this issue we will implement an incremental training method, where users will be given the possibility to classify their data, visualize and correct errors manually. As a next step we plan to offer our users the possibility to include their datasets into our training data to improve the classifier quality. As the amount of training data increases we will be able not only to provide more accurate results but to also train specialized classifiers. For example, we could provide a selection of classifiers for indoor and outdoor scenes, for different seasons and scales.

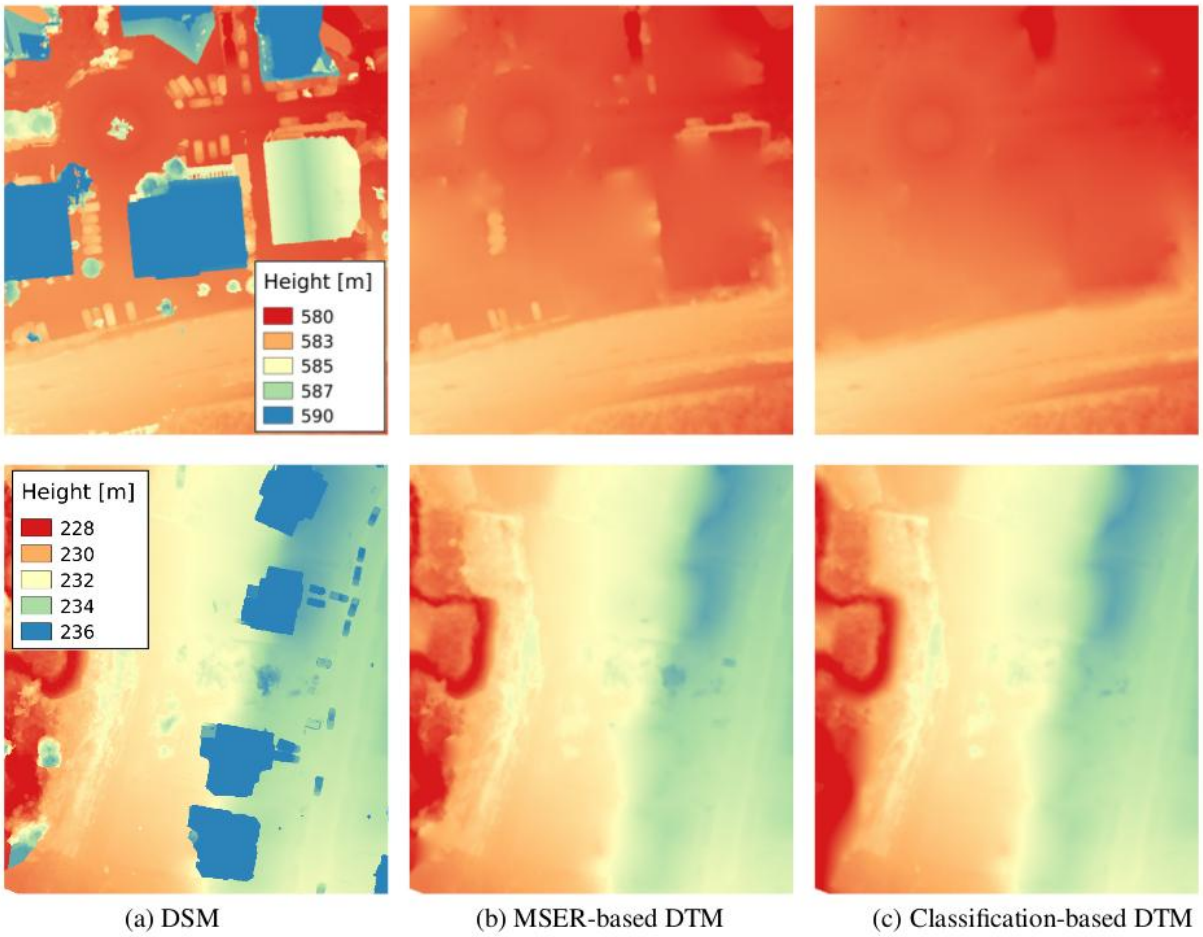

Figure 3: MSER and classification-based DTM results for the Ankeny and Building datasets. Our classification-baseapproach removes cars and low vegetation that the MSER-based

REFERENCES

Bezanson, J., Edelman, A., Karpinski, S. and Shah, V. B., 2014. Julia: A fresh approach to numerical computing. arXiv preprint arXiv:1411.1607.

Böhm, J., Bredif, M., Gierlinger, T., Krammer,¨ M., Lindenberg, R., Liu, K., Michel, F. and Sirmacek, B., 2016. The IOMULUS urban showcase: automatic tree classification and identification in huge mobile mapping point points. ISPRS - International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences XLI-B3, pp. 301–307.

Brenner, C., 2000. Towards fully automatic generation of city models. In: IAPRS, pp. 85–92.

Brodu, N. and Lague, D., 2012. 3D terrestrial LIDAR data classification of complex natural scenes using a multi-scale dimensionality criterion: Applications in geomorphology. ISPRS Journal of Photogrammetry and Remote Sensing 68, pp. 121–134.

Descombe X., Ortner M., Zerubia J., 2007. A Marked Point Process of Rectangles and Segments for Automatic Analysis of Digital Elevation Models. In: IEEE Transactions on Pattern Analysis & Machine Intelligence, vol. 30, no. , pp. 105-119.

Dorninger, P. and Nothegger, C., 2007. 3d segmentation of unstructured point clouds for building modelling. Interna-tional Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences 35(3/W49A), pp. 191–196.

eCognition, 2017. Trimble, www.ecognition.com.

Friedman, J., Hastie, T. and Tibshirani, R., 2001. The elements of statistical learning. Vol. 1, Springer series in statistics Springer, Berlin.

GlobalMapper, 2017. Blue marble geographics, www.bluemarblegeo.com.

Haala, N. and Brenner, C., 1999. Extraction of buildings and trees in urban environments. ISPRS Journal of Pho-togrammetry and Remote Sensing 54, pp. 130–137.

Haala, N., Brenner, C. and Anders, K.-H., 1998. 3d urban GIS from laser altimeter and 2D map data. International Archives of Photogrammetry and Remote Sensing 32, pp. 339–346.

Hackel, T., Wegner, J. D. and Schindler, K., 2016. Fast semantic segmentation of 3D point clouds with strongly varying density. ISPRS Annals of the Photogrammetry, Remote Sensing and Spatial Information Sciences, Prague, Czech Republic 3, pp. 177–184.

Hu, X. and Yuan, Y., 2016. Deep-learning-based classification for dtm extraction from als point cloud. Remote Sensing 8(9), pp. 730.Lafarge, F. and Mallet, C., 2012. Creating Large-Scale City Models from 3D Point Clouds: A Robust Approach with Hybrid Representation. International Journal of Computer Vision 99(1), pp. 69–85.

Liu, K. and Böhm, J., 2015. Classification of big point cloud data using cloud computing. The International Archives of Photogrammetry, Remote Sensing and Spatial Information Sciences 40(3), pp. 553.

Matas, J., Chum, O., Urban, M. and Pajdla, T., 2004. Robust wide-baseline stereo from maximally stable extremal regions. Image and vision computing.

Niemeyer, J., Rottensteiner, F. and Soergel, U., 2014. Contextual classification of LIDAR data and building object detection in urban areas. ISPRS Journal of Photogrammetry and Remote Sensing 87, pp. 152–165.

Oesau, S., Lafarge, F. and Alliez, P., 2016. Object classification via planar abstraction. ISPRS Annals of Photogram-metry, Remote Sensing and Spatial Information Sciences III-3, pp. 225–231.

Photoscan, 2017. Agisoft. www.agisoft.com.

Pix4Dmapper, 2017. Pix4D SA. www.pix4d.com.

Qi, C. R., Su, H., Mo, K. and Guibas, L. J., 2016. Pointnet: Deep learning on point sets for 3d classification and segmentation. arXiv preprint arXiv:1612.00593.

Rusu, R. B., Blodow, N. and Beetz, M., 2009. Fast Point Feature Histograms (FPFH) for 3d registration. IEEE, pp. 3212–3217.

Rusu, R. B., Blodow, N., Marton, Z., Soos, A. and Beetz, M., 2007. Towards 3D object maps for autonomous household robots. In: 2007 IEEE/RSJ International Conference on Intelligent Robots and Systems, IEEE, pp. 3191–3198.

Serna, A., Marcotegui, B., Goulette, F. and Deschaud, J.-E., 2014. Paris-rue-Madame database: a 3D mobile laser scanner dataset for benchmarking urban detection, segmentation and classification methods. In: 4th International Conference on Pattern Recognition, Applications and Methods ICPRAM 2014.

Shapovalov, R. and Velizhev, A., 2011. Cutting-Plane Training of Non-associative Markov Network for 3d Point Cloud Segmentation. IEEE, pp. 1–8.

Shu, Z., Sun, K., Qiu, K. and Ding, K., 2016. Pairwise SVM for on-board urban road LIDAR classification. ISPRS - International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences XLI-B1, pp. 109–113.

Sithole, G. and Vosselman, G., 2003. Automatic structure detection in a point-cloud of an urban landscape. In: Remote Sensing and Data Fusion over Urban Areas, 2003. 2nd GRSS/ISPRS Joint Workshop on, IEEE, pp. 67–71.

Smith, A. R., 1978. Color gamut transform pairs. ACM SIGGRAPH Computer Graphics 12(3), pp. 12–19.

Tombari, F., Salti, S. and Di Stefano, L., 2010. Unique signatures of histograms for local surface description. In European conference on computer vision, Springer, pp. 356–369.

Unger, M., Pock, T., Grabner, M., Klaus, A. and Bischof, H., 2009. A variational approach to semiautomatic generation of digital terrain models. Advances in Visual Computing pp. 1119–1130.

Vosselman, G., Dijkman, S. and others, 2001. 3D building model reconstruction from point clouds and ground plans.

International archives of photogrammetry remote sensing and spatial information sciences 34(3/W4), pp. 37–44.

Wang, Z. and Schenk, T., 2000. Building extraction and reconstruction from LIDAR data. International Archives of Photogrammetry and Remote Sensing 33(B3/2; PART 3), pp. 958–964.

Weinmann, M., Jutzi, B. and Mallet, C., 2013. Feature relevance assessment for the semantic interpretation of 3D point cloud data. ISPRS Annals of Photogrammetry, Remote Sensing and Spatial Information Sciences II-5/W2, pp. 313–318.

Weinmann, M., Schmidt, A., Mallet, C., Hinz, S., Rottensteiner, F. and Jutzi, B., 2015a. Contextual classification of point cloud data by exploiting individual 3d neighborhoods. ISPRS Annals of Photogrammetry, Remote Sensing and Spatial Information Sciences II-3/W4, pp. 271–278.

Weinmann, M., Urban, S., Hinz, S., Jutzi, B. and Mallet, C., 2015b. Distinctive 2D and 3D features for automated large-scale scene analysis in urban areas. Computers & Graphics 49, pp. 47–57.

Xiao, J. and Furukawa, Y., 2014. Reconstructing the worlds museums. International Journal of Computer Vision 110(3), pp. 243–258.

Xu, Y., Tuttas, S., Heogner, L. and Stilla, U., 2016. Classification of photogrammetric point clouds of scaffolds for con-struction site monitoring using Subspace Clustering and PCA. ISPRS - International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences XLI-B3, pp. 725–732.

Zhou, M., Li, C. R., Ma, L. and Guan, H. C., 2016. Land cover classification from full-waveform LIDAR data based on Support Vector Machines. ISPRS - International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences XLI-B3, pp. 447–452.

BIOGRAPHICAL NOTES

Carlos Becker is a part of the Machine Learning R&D Team at Pix4D and holds a PhD in Computer Vision from EPFL, Switzerland. His research interests include ML techniques applied to computer vision-based algorithms.

Elena Rosinskaya is a part of the Machine Learning R&D Team at Pix4D and holds a Master degree in Applied Mathematics and Computer Science from Lomonosov Moscow State University, Russia. Her interests include ML techniques applied and computer vision.

Nicolai Häni participated to this work while being part of the Machine Learning R&D Team at Pix4D. He is pursuing a PhD in University of Minnesota, USA. His research interests include ML techniques applied to robotics.

Emmanuel d'Angelo is leading the Machine Learning R&D Team at Pix4D and holds a PhD in Electrical Engineering from EPFL, Switzerland. His research interests include image processing and computer vision algorithms development.

Christoph Strecha is the CEO and founder of Pix4D and holds a PhD in Computer Vision from the Catholic University of Leuven, Belgium. His research interests include structure from motion techniques and city modeling.

CONTACTS

Elena Rosinskaya
Pix4D SA
Innovation Park, Building F EPFL, 1015
Lausanne
SWITZERLAND
Tel. +41 21 552 05 90
Fax +
Email: elena.rosinskaya@pix4d.com