Involving Different Neighborhood Types for the Analysis of Low-level Geometric 2D and 3D Features and their Relevance for Point Cloud Classification

MARTIN WEINMANN[1], CLÉMENT MALLET[2] & BORIS JUTZI[1]

Abstract: In this paper, we address the automatic interpretation of 3D point cloud data in terms of associating a (semantic) class label to each 3D point. In particular, we aim at analyzing the behavior of standard handcrafted low-level geometric 2D and 3D features for different neighborhood types. For this purpose, we present a framework that considers four neighborhood definitions as the basis for calculating a set of 18 low-level geometric 2D and 3D features which, in turn, are provided as input to three classifiers relying on different learning principles. We demonstrate the performance of our framework on a benchmark dataset for which a labeling with respect to three structural classes (linear, planar and volumetric structures) as well as a labeling with respect to five semantic classes (wire, pole/trunk, façade, ground and vegetation) is available. The derived results clearly reveal that the suitability of the considered neighborhood types and thus the relevance of respectively extracted features with respect to the classification task varies significantly.

1 Introduction

The automatic analysis of 3D point cloud data is a topic of major interest in photogrammetry, remote sensing, computer vision and robotics. Among a variety of tasks, attention has been paid in particular to 3D point cloud classification, where the objective consists in automatically labeling the given 3D points with respect to pre-defined classes. Thereby, the main challenges are represented by the irregular point sampling, very different types of objects, the typically high complexity of observed scenes and the high computational burden arising from the consideration of large 3D point clouds in combination with a rich diversity of available features.

Many investigations on 3D point cloud classification rely on the use of geometric features and, consequently, a respective processing pipeline typically involves (1) the recovery of a local neighborhood for each point of the 3D point cloud, (2) the extraction of geometric features via a consideration of the spatial arrangement of all points inside the local neighborhood and (3) a supervised classification via a classifier trained on representative training data. Accordingly, the quality of derived classification results relies on the selected neighborhood definition, the extracted features and the involved classifier.

In this paper, we focus on analyzing the behavior of standard handcrafted low-level geometric 2D and 3D features for different neighborhood types. For this purpose, we consider four neighborhood definitions as the basis for feature extraction: (1) a cylindrical neighborhood with a fixed radius of 1m, (2) a spherical neighborhood with a fixed radius of 1m, (3) a spherical

[1] Karlsruhe Institute of Technology (KIT), Institute of Photogrammetry and Remote Sensing, Englerstraße 7, 76131 Karlsruhe, Germany, E-Mail: [martin.weinmann, boris.jutzi]@kit.edu

[2] Université Paris-Est, LASTIG MATIS, IGN, ENSG, 73 avenue de Paris, 94160 Saint-Mandé, France, E-Mail: clement.mallet@ign.fr

neighborhood formed by the $k = 50$ nearest neighbors, and (4) a spherical neighborhood formed by the k_{opt} nearest neighbors whereby the optimal scale parameter k_{opt} is selected for each 3D point individually via *eigenentropy-based scale selection* (WEINMANN et al. 2015a). Based on these neighborhoods, we calculate a set of 18 low-level geometric 2D and 3D features and, subsequently, we evaluate their relevance for the classification task. For this purpose, we present classification results achieved for three classifiers relying on different learning principles (instance-based learning, probabilistic learning and ensemble learning) and for different classification tasks. By involving a benchmark dataset for which a labeling with respect to three structural classes (*linear structures*, *planar structures* and *volumetric structures*) as well as a labeling with respect to five semantic classes (*wire*, *pole/trunk*, *façade*, *ground* and *vegetation*) is available as shown in Fig. 1, we demonstrate that the relevance of extracted features with respect to a classification task may significantly vary depending on the selected neighborhood type.

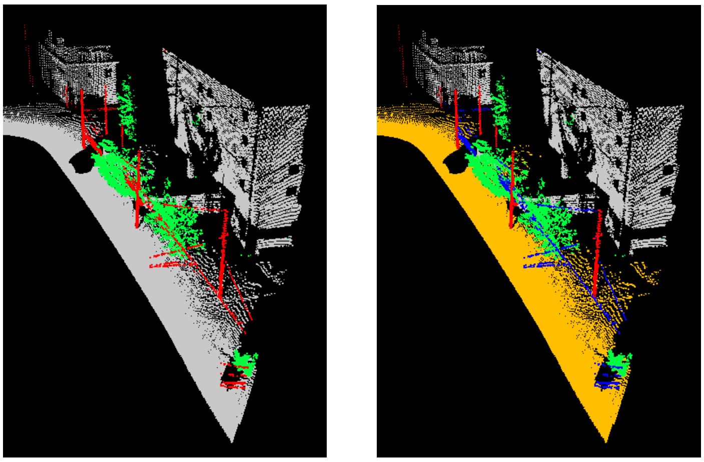

Fig. 1: Benchmark dataset with a reference labeling with respect to three structural classes (left), which indicate *linear structures* (red), *planar structures* (gray) and *volumetric structures* (green), and with a reference labeling with respect to five semantic classes (right), which are represented by *wire* (blue), *pole/trunk* (red), *façade* (gray), *ground* (brown) and *vegetation* (green).

This paper represents an extension of our previous work (WEINMANN et al. 2015a), where we presented a framework for semantic 3D point cloud interpretation based on optimal neighborhoods, relevant features and efficient classifiers. From this framework, we make use of (1) the eigenentropy-based scale selection in order to recover the optimal scale parameter k_{opt} for each 3D point individually and (2) the extraction of a set of 18 low-level geometric 2D and 3D features. In addition, we now consider further neighborhood types as a basis for extracting the same set of handcrafted features, and we evaluate their relevance with respect to two different classification tasks. Thus, our contribution consists of an extended framework which involves

- the recovery of local neighborhoods based on conceptually different neighborhood types,
- the extraction of low-level geometric 2D and 3D features on the basis of different neighborhood types, and
- a classification via standard classifiers relying on different learning principles.

After briefly summarizing related work (Section 2), we proceed with describing the proposed framework in detail (Section 3). Subsequently, we present the results derived for a benchmark dataset (Section 4), whereby we focus on both the behavior of geometric features for different neighborhood types and the classification results derived for two different classification tasks. Finally, we provide concluding remarks and suggestions for future work (Section 5).

2 Related Work

The semantic classification of 3D point cloud data typically relies on the use of geometric features. The extraction of such geometric features often exploits characteristics of the local 3D structure, i.e. local neighborhoods have to be recovered (Section 2.1). On the basis of the 3D points within the recovered local neighborhoods, a variety of geometric features can be calculated (Section 2.2). These features are then provided as input to the classifier that is trained on representative training data (Section 2.3).

2.1 Neighborhood Recovery

In general, there are very different strategies for recovering local neighborhoods for the points of a 3D point cloud. Among the most commonly used neighborhood types, one can find spherical neighborhoods parameterized by a radius (LEE & SCHENK 2002), cylindrical neighborhoods parameterized by a radius (FILIN & PFEIFER 2005), spherical neighborhoods parameterized by the number of nearest neighbors with respect to the Euclidean distance in 3D space (LINSEN & PRAUTZSCH 2001) and cylindrical neighborhoods parameterized by the number of nearest neighbors with respect to the Euclidean distance in 2D space (NIEMEYER et al. 2014). However, all these neighborhood types rely on a scale parameter which is represented by either a radius or the number of nearest neighbors. To define this scale parameter, prior knowledge about the scene and/or the data is typically involved. Furthermore, the same value of the scale parameter is commonly selected for all points of the 3D point cloud. To allow for more variability and thus account for the fact that structures of different classes might favor a different neighborhood size, data-driven approaches for optimal neighborhood size selection have been presented (MITRA & NGUYEN 2003; LALONDE et al. 2005; DEMANTKÉ et al. 2011; WEINMANN et al. 2015a).
Instead of describing the local 3D structure at a specific scale, it is possible to describe the local 3D structure at different scales and thus also how the local 3D geometry behaves across scales. Respective approaches for instance focus on the combination of multiple spherical neighborhoods with different radii (BRODU & LAGUE 2012) or the combination of multiple cylindrical neighborhoods with different radii (NIEMEYER et al. 2014; SCHMIDT et al. 2014). The combination of different entities such as voxels, blocks and pillars results in multi-type neighborhoods (HU et al. 2013) and the use of multi-scale, multi-type neighborhoods in the form of a combination of both spherical and cylindrical neighborhoods with different scale parameter has been proposed in (BLOMLEY et al. 2016).

2.2 Feature Extraction

Based on the recovered local neighborhoods, a variety of geometric features can be calculated. Among these, the eigenvalues of the 3D structure tensor (JUTZI & GROSS 2009) are often used to derive the *local 3D shape features* presented in (WEST et al. 2004; PAULY et al. 2003) which are widely used for classifying 3D point cloud data. To address further characteristics of the given 3D point cloud data, it has been proposed to additionally extract features based on angular characteristics (MUNOZ et al. 2009), height and plane characteristics (MALLET et al. 2011), a variety of low-level geometric 3D and 2D features (WEINMANN et al. 2015a), moments and height features (HACKEL et al. 2016), or specific descriptors addressing surface properties, slope, height characteristics, vertical profiles and 2D projections (GUO et al. 2015).

2.3 Classification

The calculated features are provided as input to the classifier which may rely on different strategies. On the one hand, it is possible to involve standard classifiers such as a Random Forest classifier (CHEHATA et al. 2009) or a Support Vector Machine classifier (MALLET et al. 2011). A comprehensive comparison of a variety of standard classifiers relying on different learning principles is presented in (WEINMANN et al. 2015a) and reveals that a Random Forest classifier provides a good trade-off between classification accuracy and computational efficiency. On the other hand, it is possible to take into account that the derived labeling should be spatially regular which can be enforced via statistical models of context such as Associative Markov Networks (MUNOZ et al. 2009), non-Associative Markov Networks (SHAPOVALOV et al. 2010), or Conditional Random Fields (NIEMEYER et al. 2014; WEINMANN et al. 2015b).

3 Methodology

In the scope of this paper, we focus on a framework for point cloud classification, which allows statements about the suitability of involved neighborhood types and the relevance of respectively extracted features with respect to the considered classification task. This framework is illustrated in Fig. 2 and consists of three components addressing neighborhood recovery (Section 3.1), feature extraction (Section 3.2) and classification (Section 3.3).

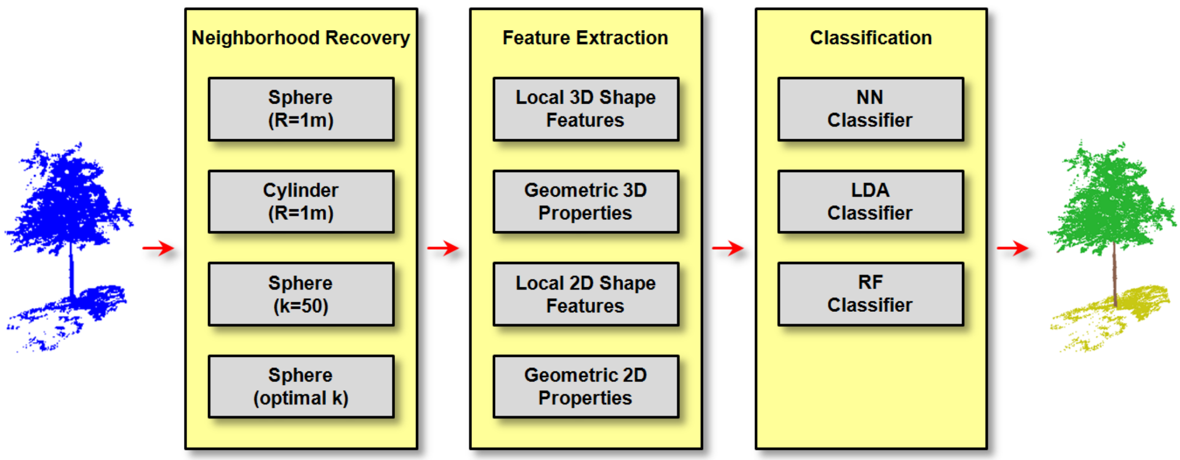

Fig. 2: The proposed framework for point cloud classification with four options for neighborhood recovery, four types of low-level geometric features and three options for supervised classification (NN: Nearest Neighbor; LDA: Linear Discriminant Analysis; RF: Random Forest).

3.1 Neighborhood Recovery

As geometric features describing a point $\mathbf{X}_0$ are typically derived by considering the local surrounding of $\mathbf{X}_0$, a crucial prerequisite for calculating geometric features consists in the choice of an appropriate neighborhood definition. We therefore include four different neighborhood definitions in our framework:

- a spherical neighborhood $N_{s,1\mathrm{m}}$, where a sphere is centered at $\mathbf{X}_0$ and has a radius of 1m,
- a cylindrical neighborhood $N_{c,1\mathrm{m}}$, where a cylinder is centered at $\mathbf{X}_0$, has a radius of 1m and is oriented along the vertical direction,
- a spherical neighborhood $N_{k=50}$ comprising the $k = 50$ nearest neighbors of $\mathbf{X}_0$ with respect to the Euclidean distance in 3D space, and
- a spherical neighborhood $N_{k,\mathrm{opt}}$ comprising the optimal number k_{opt} of nearest neighbors of $\mathbf{X}_0$ with respect to the Euclidean distance in 3D space, whereby k_{opt} is selected for each point individually via *eigenentropy-based scale selection* (WEINMANN et al. 2015a).

3.2 Feature Extraction

Once a local neighborhood has been derived for each point $\mathbf{X}_0$, the spatial arrangement of points within that local neighborhood can be described via geometric features. Here, we focus on the use of low-level geometric features that have been presented in (WEINMANN et al. 2015a). These features comprise the *local 3D shape features* of linearity, planarity, sphericity, omnivariance, anisotropy, eigenentropy, sum of eigenvalues and change of curvature (WEST et al. 2004; PAULY et al. 2003) that rely on the eigenvalues of the 3D structure tensor. In addition, features are defined based on *geometric 3D properties* that are represented by the height of $\mathbf{X}_0$, the distance between $\mathbf{X}_0$ and the farthest point in the local neighborhood, the local point density, the verticality, and the maximum difference as well as the standard deviation of the height values corresponding to those points within the local neighborhood (WEINMANN et al. 2015a). As the presence of different man-made objects can be assumed in urban scenes, some further features rely on a 2D projection of all points within the local neighborhood onto a horizontally oriented plane. Using the 2D projections, it is possible to define *local 2D shape features* in terms of the sum and the ratio of the eigenvalues of the 2D structure tensor (WEINMANN et al. 2015a), and *geometric 2D properties* are given by the distance between the projection of $\mathbf{X}_0$ and the farthest point in the local 2D neighborhood and the local point density in 2D space.

3.3 Supervised Classification

For classification, we integrate three classifiers relying on different learning principles into our framework. This allows us to draw more general conclusions about the suitability of different neighborhood types and thus the relevance of respectively extracted features with respect to the classification task:

(1) *Nearest Neighbor* (*NN*) *classifier*:
This classifier relies on the principle of instance-based learning, where the induction is delayed to the prediction stage. Accordingly, there is no training stage. In the prediction stage, each feature vector belonging to a sample in the test set is compared to the feature

vectors corresponding to all samples in the training data and the class label associated with the most similar feature vector in the training data is selected. Due to the typically large number of required comparisons of feature vectors (here with respect to the Euclidean distance), the prediction stage is likely to cause a higher computational burden.

(2) *Linear Discriminant Analysis* (*LDA*) *classifier*:
This classifier relies on the principle of probabilistic learning. In the training stage, a multivariate Gaussian distribution is fitted to the training data, i.e. the parameters of a Gaussian distribution have to be estimated for each class, assuming that the instances of each class follow a Gaussian distribution in the feature space. Thereby, the LDA classifier involves the simplifying assumption that the same covariance matrix can be used for all classes so that only their means vary. In the prediction stage, it is evaluated with which probability each feature vector belonging to a sample in the test set is assigned to the different classes, and the class with the maximum probability is selected.

(3) *Random Forest* (*RF*) *classifier*:
This classifier relies on the principle of ensemble learning (BREIMAN 2001). In the training stage, the strategy of bagging (BREIMAN 1996) is applied to generate an ensemble of randomly trained decision trees. More specifically, different subsets of the training data are randomly drawn with replacement so that an individual decision tree can be trained for each subset. In the prediction stage, for each feature vector belonging to a sample in the test set, each of the randomly trained decision tree casts a vote for one of the defined classes and, typically, the majority vote is selected to obtain an appropriate prediction of the respective class label.

4 Experimental Results

In the following, we first present the involved benchmark dataset (Section 4.1). Subsequently, we focus on the behavior of geometric features for different neighborhood types and present the derived classification results (Section 4.2). Finally, we discuss the derived results with respect to different aspects (Section 4.3).

4.1 Dataset

We evaluate the performance of our framework on the *Oakland 3D Point Cloud Dataset* (MUNOZ et al. 2009), a widely used benchmark dataset acquired with a mobile laser scanning system in the vicinity of the CMU campus in Oakland, USA. The dataset is split into a training set comprising about 36.9k points, a validation set comprising about 91.5k points and a test set comprising about 1.3M points. For each point, a reference labeling with respect to three structural classes (*linear structures*, *planar structures* and *volumetric structures*) is available as well as a reference labeling with respect to five semantic classes (*wire*, *pole*/*trunk*, *façade*, *ground* and *vegetation*) as shown in Fig. 1. To distinguish between both classification tasks, we refer to *Oakland-3C* and *Oakland-5C*, respectively.

4.2 Results

First, we focus on the structural behavior of the low-level geometric features for the different neighborhood types. Thereby, we exemplarily consider the behavior of the three dimensionality features of linearity, planarity and sphericity which is illustrated in Fig. 3. It can be observed that the different neighborhood types have a significantly different impact on feature extraction.

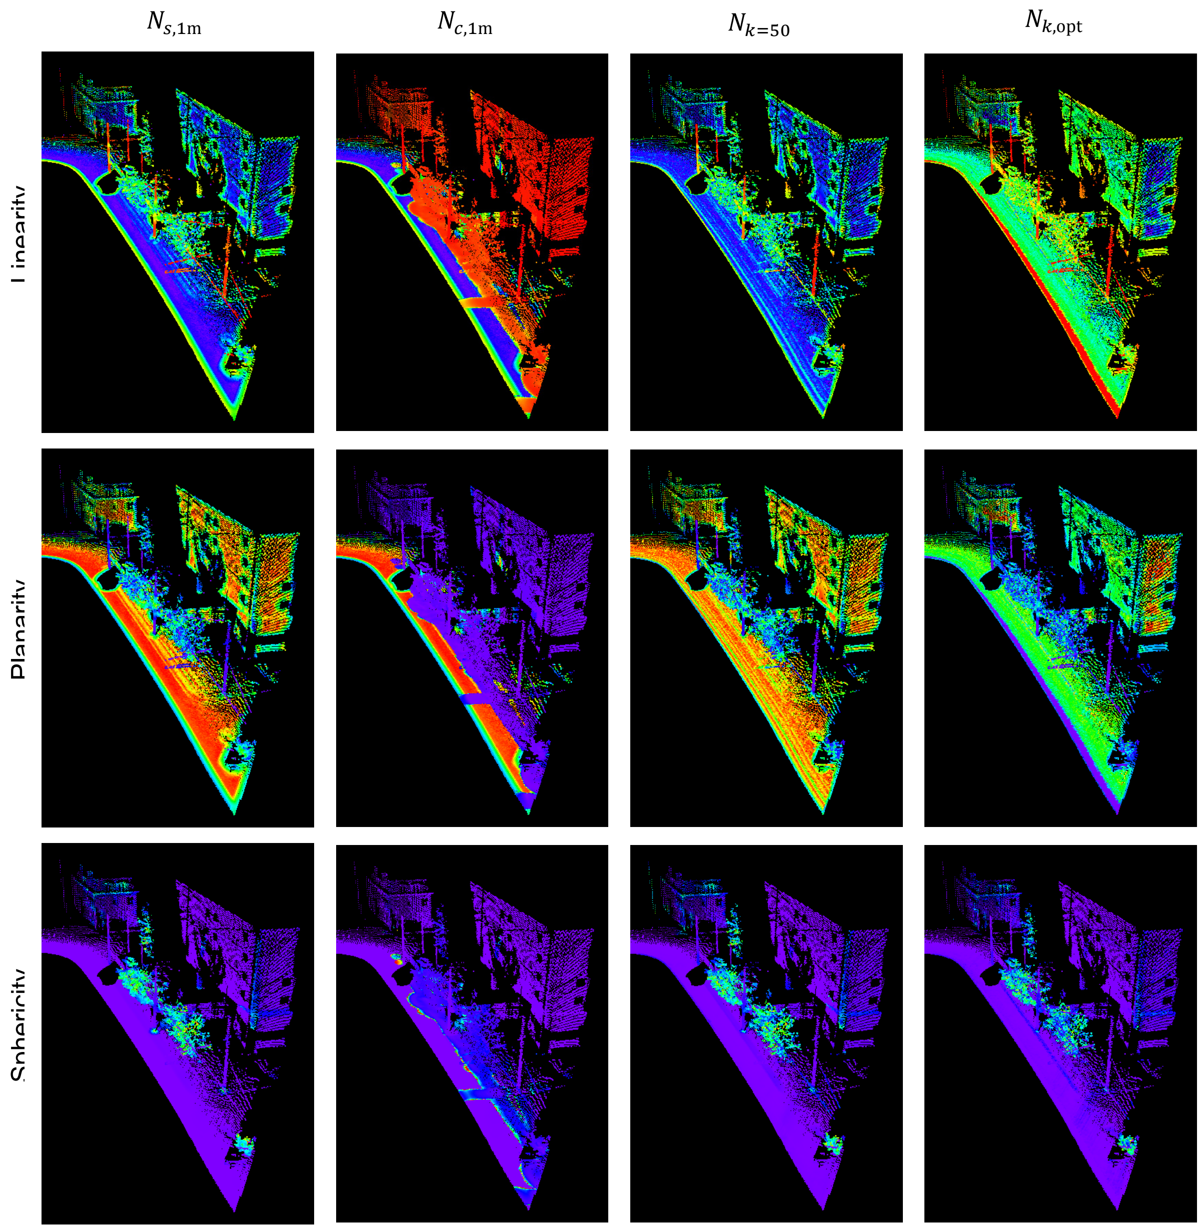

Fig. 3: Behavior of the dimensionality features of linearity, planarity and sphericity for the considered scene: the color encoding indicates high values close to 1 in red and reaches via yellow, green, cyan and blue to violet for low values close to 0.

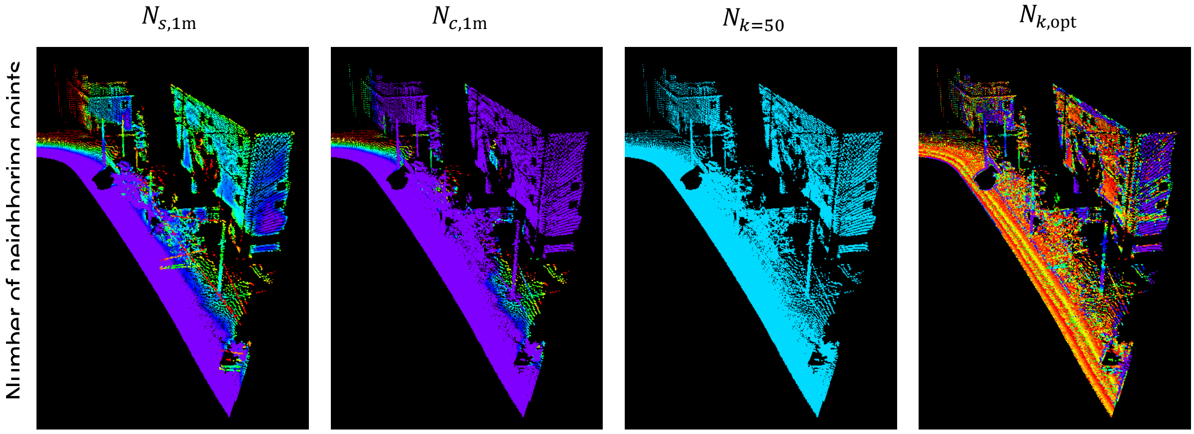

Fig. 4: Number of neighboring points for the different neighborhood types: the color encoding indicates neighborhoods with 10 or less points in red and reaches via yellow, green, cyan and blue to violet for 100 and more points.

Additionally, we consider the number of neighbors within the local neighborhood which is visualized in Fig. 4. Due to their significantly different impact on feature extraction, it can be expected that the different neighborhood types might also significantly differ in their suitability with respect to the classification task. To verify this, we use the different neighborhood types as the basis for extracting the 18 low-level geometric 2D and 3D features which, in turn, are provided as input to three classifiers relying on different learning principles. In the training stage, we take into account that an unbalanced distribution of training examples per class might have a detrimental effect on the classification results (CRIMINISI & SHOTTON 2013) and hence use a reduced training set comprising 1000 training examples per class. In the prediction stage, we estimate the class labels for the validation set. For *Oakland-3C*, the derived classification results are provided in Tab. 1 and visualized in Fig. 5. The quantitative results reveal an overall accuracy in the range of about 67-94% and a κ-value in the range of about 28-76%. For *Oakland-5C*, the derived classification results are provided in Tab. 2 and visualized in Fig. 6. Here, the quantitative results reveal an overall accuracy in the range of about 68-96%, whereas the κ-value is in the range of about 49-90%.

4.3 Discussion

The derived results clearly reveal that the spherical and cylindrical neighborhood types parameterized by a radius result in local neighborhoods that comprise a larger number of 3D points, while the neighborhoods recovered via eigenentropy-based scale selection tend to be comparably small (Fig. 4). Furthermore, it can be observed that the selected cylindrical neighborhood type is not well-suited for classifying mobile laser scanning data as extracted features exhibit a rather non-intuitive behavior with seam effects (Fig. 3) and do hence not lead to appropriate classification results (Tab. 1, Tab. 2, Fig. 5 and Fig. 6). In contrast, the use of the other neighborhood types results in adequate classification results for almost all cases, with one exception for the combination of using locally optimal neighborhoods and a RF classifier. For that case, misclassifications mainly occur in regions of low point density (Fig. 5 and Fig. 6).

Tab. 1: Classification results obtained for *Oakland-3C* when using four different neighborhood types.

Neighborhood Type	Overall Accuracy [%]			κ [%]		
	NN	LDA	RF	NN	LDA	RF
$N_{s,1\mathrm{m}}$	90.80	92.78	91.74	66.09	72.30	69.47
$N_{c,1\mathrm{m}}$	84.00	78.29	82.26	50.57	41.55	48.30
$N_{k=50}$	92.41	91.27	92.52	70.49	67.60	71.39
$N_{k,\mathrm{opt}}$	67.86	90.91	93.83	28.33	67.47	75.24

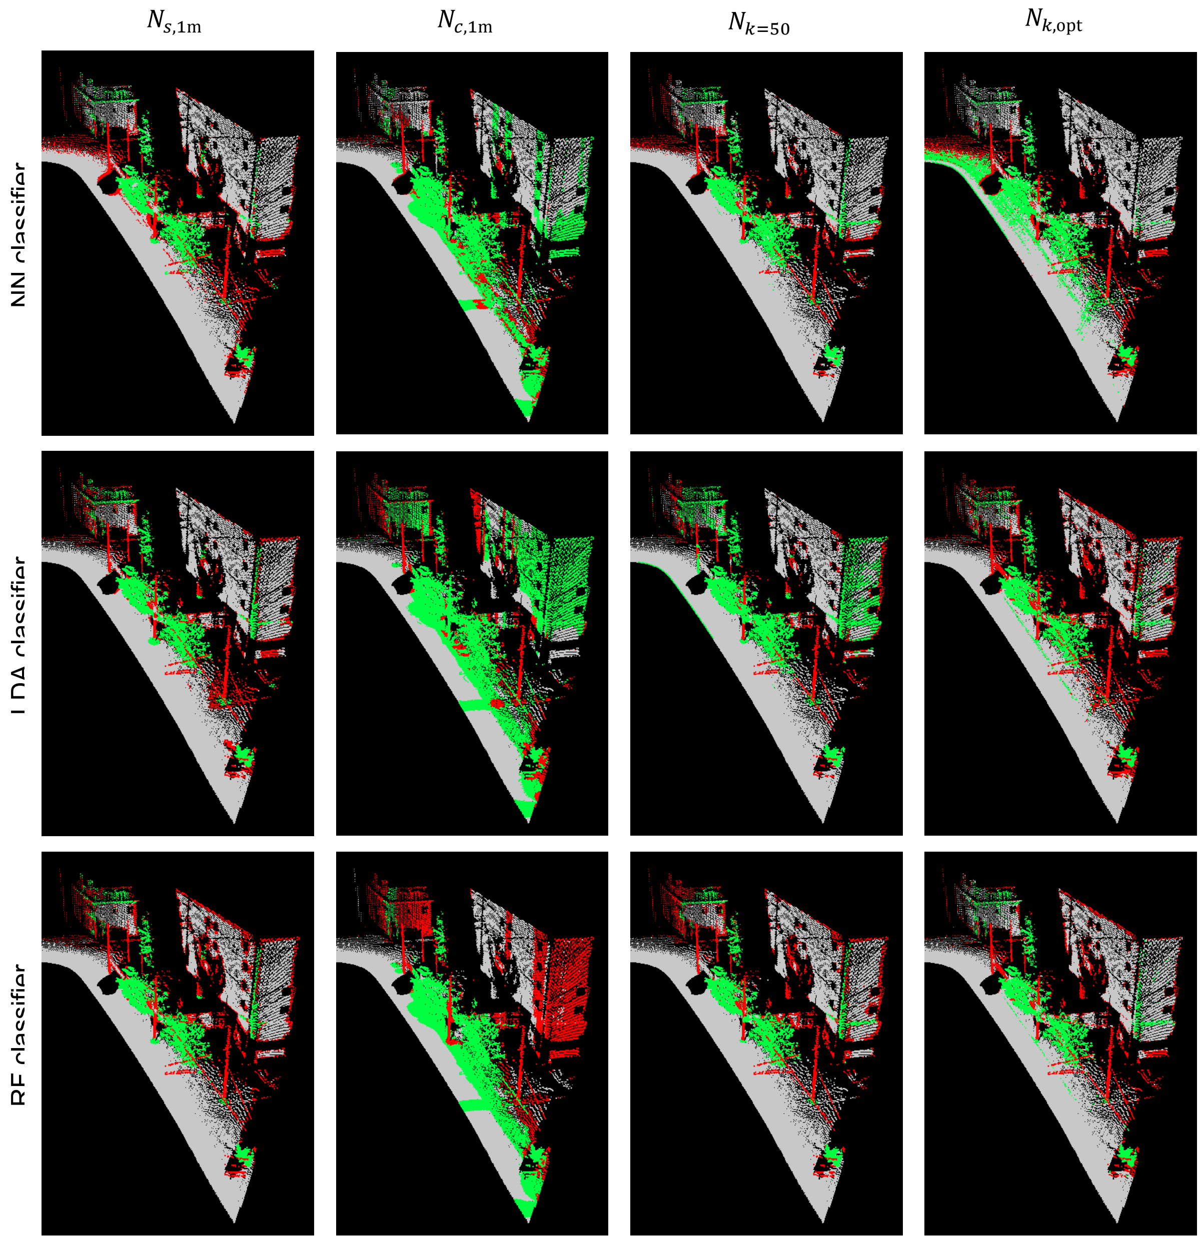

Fig. 5: The observed scene classified with respect to three structural classes which are represented by *linear structures* (red), *planar structures* (gray) and *volumetric structures* (green).

Tab. 2: Classification results obtained for *Oakland-5C* when using four different neighborhood types.

Neighborhood Type	Overall Accuracy [%]			κ [%]		
	NN	LDA	RF	NN	LDA	RF
$N_{s,1\mathrm{m}}$	90.26	95.11	93.73	79.18	88.78	85.71
$N_{c,1\mathrm{m}}$	85.22	86.73	85.89	70.01	72.68	71.18
$N_{k=50}$	91.09	95.28	95.36	80.61	89.06	89.26
$N_{k,\mathrm{opt}}$	68.06	94.38	94.15	49.17	87.30	86.66

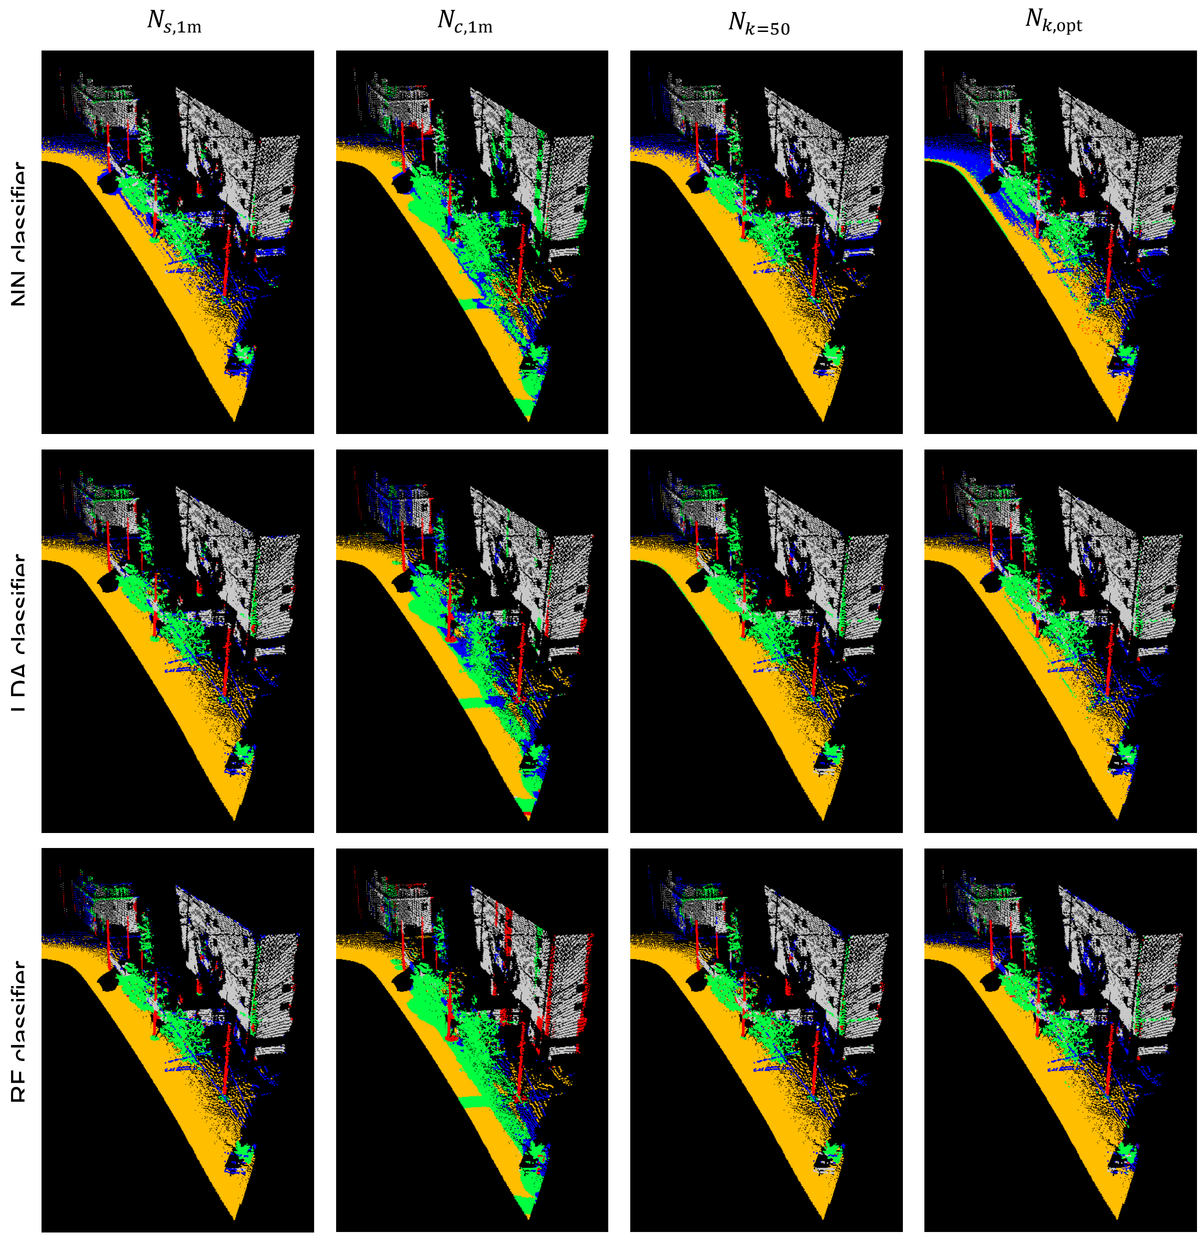

Fig. 6: The observed scene classified with respect to five semantic classes which are represented by *wire* (blue), *pole/trunk* (red), *façade* (gray), *ground* (brown) and *vegetation* (green).

In addition, it should be taken into account that the neighborhood types $N_{s,1\mathrm{m}}$, $N_{c,1\mathrm{m}}$ and $N_{k=50}$ rely on a scale parameter that is selected heuristically and kept identical for all points of the point cloud. For a 3D point cloud acquired with a different scanning device or within a different scene, the point sampling might exhibit different characteristics so that the selected scale parameters might not be that suitable anymore. In contrast, the neighborhood type $N_{k,\mathrm{opt}}$ allows a data-driven selection of the scale parameter for each individual 3D point. Accordingly, it takes into account that different classes might favor a different neighborhood size. Furthermore, it better accounts for variations in point density (Fig. 5 and Fig. 6), particularly when using a LDA classifier or a RF classifier.

Among the considered classifiers, the LDA classifier and the RF classifier yield results of similar quality and tend to outperform the NN classifier for almost all cases. Due to the simplifying assumptions made with the LDA classifier – which cannot be guaranteed for the acquired data – using the RF classifier seems to be favorable.

Consequently, we conclude that local neighborhoods selected via *eigenentropy-based scale selection* are to be preferred as the basis for feature extraction, as respectively extracted features allow deriving appropriate classification results when using a RF classifier which we consider as the favorable option for classification.

5 Conclusions

In this paper, we have considered different neighborhood types and analyzed the respective behavior of low-level geometric 2D and 3D features and their relevance for point cloud classification. We have presented a framework allowing for the use of four neighborhood types as the basis for extracting a set of 18 low-level geometric 2D and 3D features which, in turn, are provided as input to classification. For the latter, we involved classifiers relying on three different learning principles. The derived results for a benchmark dataset with two reference labelings indicate that the relevance of extracted features with respect to a classification task may significantly vary depending on the selected neighborhood type. While the selected cylindrical neighborhood type clearly reveals less suitability with respect to the classification task, particularly the neighborhood types relying on a number of nearest neighbors lead to appropriate classification results, especially for the stronger classifiers based on probabilistic and ensemble learning.

In the scope of future work, it seems worth investigating the behavior of multi-scale neighborhoods in more detail. Thereby, the computational burden is typically significantly higher so that efficient approaches are required as for instance a recently proposed approach relying on a scale pyramid that is created by repeatedly downsampling a given 3D point cloud (HACKEL et al. 2016). Additionally, we plan to address the (partially) noisy behavior of the derived labeling. In this regard, we intend to use different approaches to impose spatial regularity on the classification results. A further aim consists in the transfer of the derived classification results to the level of single objects in the observed scene as a first step towards scene understanding.

6 References

BLOMLEY, R., JUTZI, B. & WEINMANN, M., 2016: 3D semantic labeling of ALS point clouds by exploiting multi-scale, multi-type neighborhoods for feature extraction. Proceedings of the International Conference on Geographic Object-Based Image Analysis, 1-8.

BREIMAN, L., 1996: Bagging predictors. Machine Learning **24**(2), 123-140.

BREIMAN, L., 2001: Random forests. Machine Learning **45**(1), 5-32.

BRODU, N. & 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 & Remote Sensing **68**, 121-134.

CHEHATA, N., GUO, L. & MALLET, C., 2009: Airborne lidar feature selection for urban classification using random forests. The International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences **38**(3/W8), 207-212.

CRIMINISI, A. & SHOTTON, J., 2013: Decision forests for computer vision and medical image analysis. Springer, London, UK.

DEMANTKÉ, J., MALLET, C., DAVID, N. & VALLET, B., 2011: Dimensionality based scale selection in 3D lidar point clouds. The International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences **38**(5/W12), 97-102.

FILIN, S. & PFEIFER, N., 2005: Neighborhood systems for airborne laser data. Photogrammetric Engineering & Remote Sensing **71**(6), 743-755.

GUO, B., HUANG, X., ZHANG, F. & SOHN, G., 2015: Classification of airborne laser scanning data using JointBoost. ISPRS Journal of Photogrammetry & Remote Sensing **100**, 71-83.

HACKEL, T., WEGNER, J. D. & 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 **III** (3), 177-184.

HU, H., MUNOZ, D., BAGNELL, J. A. & HEBERT, M., 2013: Efficient 3-D scene analysis from streaming data. Proceedings of the IEEE International Conference on Robotics and Automation, 2297-2304.

JUTZI, B. & GROSS, H., 2009: Nearest neighbour classification on laser point clouds to gain object structures from buildings. The International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences **38**(1-4-7/W5), 1-6.

LALONDE, J.-F., UNNIKRISHNAN, R., VANDAPEL, N. & HEBERT, M., 2005: Scale selection for classification of point-sampled 3D surfaces. Proceedings of the International Conference on 3-D Digital Imaging and Modeling, 285-292.

LEE, I. & SCHENK, T., 2002: Perceptual organization of 3D surface points. The International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences, **34**(3A), 193-198.

LINSEN, L. & PRAUTZSCH, H., 2005: Local versus global triangulations. Proceedings of Eurographics, 257-263.

MALLET, C., BRETAR, F., ROUX, M., SOERGEL, U. & HEIPKE, C., 2011: Relevance assessment of full-waveform lidar data for urban area classification. ISPRS Journal of Photogrammetry & Remote Sensing **66**(6), S71-S84.

MITRA, N. J. & NGUYEN, A., 2003: Estimating surface normals in noisy point cloud data. Proceedings of the Annual Symposium on Computational Geometry, 322-328.

MUNOZ, D., BAGNELL, J. A., VANDAPEL, N. & HEBERT, M., 2009: Contextual classification with functional max-margin Markov networks. Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, 975-982.

NIEMEYER, J., ROTTENSTEINER, F. & SOERGEL, U., 2014: Contextual classification of lidar data and building object detection in urban areas. ISPRS Journal of Photogrammetry & Remote Sensing **87**, 152-165.

PAULY, M., KEISER, R. & GROSS, M., 2003: Multi-scale feature extraction on point-sampled surfaces. Computer Graphics Forum **22**(3), 81-89.

SCHMIDT, A., NIEMEYER, J., ROTTENSTEINER, F. & SOERGEL, U., 2014: Contextual classification of full waveform lidar data in the Wadden Sea. IEEE Geoscience and Remote Sensing Letters **11**(9), 1614-1618.

SHAPOVALOV, R., VELIZHEV, A. & BARINOVA, O., 2010: Non-associative Markov networks for 3D point cloud classification. The International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences **38**(3A), 103-108.

WEINMANN, M., JUTZI, B., HINZ, S. & MALLET, C., 2015a: Semantic point cloud interpretation based on optimal neighborhoods, relevant features and efficient classifiers. ISPRS Journal of Photogrammetry and Remote Sensing **105**, 286-304.

WEINMANN, M., SCHMIDT, A., MALLET, C., HINZ, S., ROTTENSTEINER, F. & JUTZI, B., 2015b: Contextual classification of point cloud data by exploiting individual 3D neighborhoods. ISPRS Annals of the Photogrammetry, Remote Sensing and Spatial Information Sciences **II**(3/W4), 271-278.

WEST, K. F., WEBB, B. N., LERSCH, J. R., POTHIER, S., TRISCARI, J. M. & IVERSON, A. E., 2004: Context-driven automated target detection in 3-D data. Proceedings of SPIE **5426**, 133-143.