Hyperspectral 3D Point Cloud Segmentation using RandLA-Net

Isaak Mitschke¹, Thomas Wiemann^{1,2}, Felix Igelbrink¹, and Joachim Hertzberg^{1,2}

 Osnabrück University, Institute of Computer Science Berghoffstr. 11, 49090 Osnabrück, Germany
imitschke|twiemann|figelbrink@uni-osnabrueck.de
DFKI Niedersachsen Lab, Plan Based Robot Control Berghoffstr. 11, 49090 Osnabrück, Germany
thomas.wiemann|joachim.hertzberg@dfki.de

Abstract Point clouds are commonly used in robotics to represent 3D maps. To gain further understanding of their content, it is useful to annotate such maps semantically. To segment 3D point clouds with RGB values, different solutions exist. In machine learning, pre-trained classifiers are used for this purpose. Since it is not always possible to differentiate between entities relying solely on RGB information, hyperspectral histograms can augment the 3D data. The aim of this work is to evaluate, if hyperspectral information can improve the segmentation results for ambiguous objects, e.g., streets, sidewalks, and cars using established deep learning methods. Given the reported performance on geometrical 3D data and the possibility to directly integrate point annotations, we extended the neural network RandLA-Net. In addition to the evaluation of RandLA-Net in this context, we also provide a reference dataset consisting of semantically annotated hyperspectral 3D point clouds.

Keywords: Point Clouds, Semantic Segmentation, Deep Learning, Hyperspectral Imaging

1 Introduction

Semantic maps are essential for autonomous robots to perform goal-directed and unsupervised tasks. While creating such maps, instances of known classes of objects have to be labeled in the recorded 3D environment model, which is typically a 3D point cloud. In analogy to classical image analysis, it is desirable to learn classifiers that automatically segment point clouds into previously defined classes. In recent years, advances in deep learning have led to the development of an increasing number of methods for segmenting point clouds with the help of neural networks [1]. Currently, one of the best performing neural networks for classifying 3D point clouds is RandLA-Net [2]. With RandLA-Net, it is possible to segment large-scale high resolution point clouds, as shown in experiments on the Semantic3D datasets [3]. In addition to 3D coordinates, some scanning systems measure additional surface properties like color or reflectance [1]. In [4] a terrestrial 3D scanning system is described that captures hyperspectral images and maps the recorded spectra onto the recorded 3D points. Hyperspectral imaging has already been successfully applied the field of remote sensing to analyze aerial images. In this discipline, the hyperspectral information is used to accurately distinguish between different materials [5]. Deep learning using neuronal networks can further increase the segmentation precision [6]. With the development of hyperspectral laser scanners, it has become feasible to combine both modalities. However, the combination of hyperspectral data and 3D point clouds is rarely found in current research and the classification of such data is even less common. In this paper, we present a first step to apply deep learning to segment terrestrial hyperspectral point clouds. Given the promising results of RandLA-Net, our approach builds on top of that network structure. The novelty of this work is to combine the spatial information encoded in the point clouds with the additional information in the hyperspectral images. For that, we project the hyperspectral information directly on the point clouds rather than doing the classification separately on both domains. Because no benchmark datasets for classification on hyperspectral 3D point clouds are publicly available, we also provide a labeled reference dataset. Our experiments show that the classification quality on the combined data improves significantly compared to the existing RandLA-Net classifier.

2 Related Work

Classification and segmentation of 2D images has been successfully applied in many applications like driver assistance systems, medicine, and remote sensing [7]. Many of these image classification approaches rely on deep learning using neural networks. Similar to 2D images, the objective of classifying or segmenting point clouds is to assign a label to each 3D point in a point cloud. Labels of interest may be object classes like Cars and Trees, but can also coarser specifications in terms of object categories such as Vehicle or Vegetation. Using deep learning to analyze point clouds is a relatively new field of research. Over the last years, many promising approaches to solve this task have been proposed. A milestone in research on deep learning-based approaches was PointNet [8], presented in 2017, and its extension PointNet++ [9], which was published shortly afterwards. Since then, many other methods have been developed and an increasing number of benchmark datasets for the segmentation of point clouds have been made available to compare the results of different segmentation methods. Examples for such benchmark datasets are Semantic3D [3], S3DIS [10], Paris-Lille-3D [11] and SemanticKitti [12].

Guo et al. [1] evaluated and compared the performance of the most recent point cloud segmentation methods. They reported that the best results on almost all datasets were achieved by using RandLA-Net[2]. Similar to PointNet and PointNet++, RandLA-Net also uses a point-based approach, where the raw, unorganized point cloud is fed into the network. RandLA-Net is based on an encoder/decoder architecture and takes n points with corresponding features as input. The encoder layer consists of dilated residual blocks followed by randomized downsampling. The dilated residual blocks are used to compute features taking the neighborhood of each point into account. Attentive pooling is then used to generate new aggregated features.

Hyperspectral imaging is the process of capturing digital images, where during measurement, the intensity of multiple wavelengths of light is sensed rather than three color channels (red, green, and blue) as in common digital imaging. Light visible to humans is physically electromagnetic radiation between 300 nm to 750 nm. By analyzing the spectral signature of targets, detailed assumptions about the chemical nature of the material may be inferred. The spectral signature, for example, can be utilized to differentiate between vegetation and nonvegetation. This is due to the fact that chlorophyll in vegetation absorbs light in the red and blue spectral bands in particular. In the near infrared spectrum, for example, vegetation largely reflects the light [13]. A simple way of classifying vegetation and non-vegetation based on this knowledge is to compute the Normalized Difference Vegetation Index (NDVI). Vegetation indices can also be used to determine the health of vegetation. This is used in earth observation to detect drought or disease [14]. Hyperspectral imaging as a tool for non-invasive analysis of objects is used in many domains. Examples are agriculture, food processing, astronomy, geology and environmental research [15][16][17][18]. One potential limitation of using hyperspectral data is the presence of noise in the data. The main sources for that are the image sensors themselves, which induce noise, as well as physical influences like illumination fluctuations and atmospheric effects [19]. The occurrence of noise negatively influences the performance when classifying hyperspectral images [20]. Therefore, smoothing techniques are commonly applied to enhance the recorded data. Popular examples are median, moving average or Savitzky-Golay [21] filters. Another commonly used noise reduction technique is the Minimum Noise Fraction Transform (MNF), also sometimes referred to as Noise-Adjusted Principal Components Transform (NAPC). Unlike PCA, where the principal components maximize the variance, the MNF transform minimizes the noise content. The first component has the highest signal-to-noise ratio and the last component the lowest. The MNF transform was first introduced in 1988 by Green et al. [22] and in 1990 redesigned by Lee et al. [23] and renamed to NAPC.

Another problem when working with hyperspectral data are illumination differences, resulting in different spectral intensities within the same material class. This is caused by different physical properties of materials and the surface structure of the observed materials. Especially in the near infrared range [24], scattering effects occur, which lead to a high variance in different samples of one class. To minimize this variance, the Standard Normal Variant Transform (SNV) [25] or Multiplical Scatter Correction (MSC) [26] can be applied to equalize the data.

So far, there are only few approaches that use hyperspectral data combined with spatial data. Buckley et al. [27] use hyperspectral point clouds to classify geological materials in quarries. However, the classification is performed only on the hyperspectral images. The resulting segmentation is then used as texture for the 3D scene. Other approaches originate from the field of remote sensing, where aerial hyperspectral imagery is fused with airborne laser scans. In contrast to terrestrial laserscans, remote sensing data usually has a low spatial resolution. Some papers already use such sparse 3D point clouds to segment urban [28][29] or agricultural scenes [30]. But to our knowledge, no high resolution reference datasets exist, that allow to investigate the combination of large scale 3D point cloud segmentation and hyperspectral data in the context of deep learning.

3 The Hyperspectral Semantic Street Scene Dataset

In order to annotate point clouds with hyperspectral data, the camera and the laser scanner have to be co-calibrated. The extrinsic calibration, which describes the transformation from laser scanner to camera, is used to convert the 3D world coordinates into camera coordinates. Using the intrinsic calibration of the camera, the points are projected into pixel coordinates. Thus, a corresponding color value or, in the case of hyperspectral cameras, a corresponding spectrum can be assigned to each 3D point. Igelbrink et al. [4] developed a method for an ad-hoc calibration of a terrestrial 3D laser scanner and a hyperspectral camera. The calibration procedure is performed using an automated, marker-less method instead of using reference patterns. Igelbrink's method calculates the Normalized Mutual Information (NMI), which is a measure of the similarity between two images. A panorama is generated from the point intensity values of the 3D laser scan, using the cylindrical camera model, which is then compared with the hyperspectral panorama. With a perfect calibration, the NMI value is equal to 1. The Normalized Mutual Information metric is specifically well-suited for comparing multi-modal images, such as in this case an intensity image from the laser scanner and the hyperspectral image. An optimization procedure that maximizes the NMI is then utilized to find the best parameters for the calibration. The evaluation of the procedure shows that the resulting calibration is very accurate and is generated in a few seconds through GPU acceleration.

In order to train classifiers, a lot of training data is required. For all classes, numerous examples should be available, such that the classifiers have a high level of transferability to previously unseen data [31]. In this work, we chose an urban environment as a show-case to apply deep learning on the combined data for several reasons. First, urban environments are an obvious application scenario for mobile systems. Second, they are often challenging for classical image-based segmentation, as many elements like streets and sidewalks have similar colors. To overcome these ambiguities, the inclusion of additional information, e.g., spatial relations from the 3D point clouds is necessary. Third, we expect that the mixture of different materials like concrete, asphalt and vegetation will lead to significantly different spectral signatures, that will demonstrate the benefit of including the spectral domain into the classification process. Given these prop-

erties of the input data, a trained classifier should be able to differentiate between streets, sidewalks, cars, vegetation and buildings.

To verify these assumptions, we created a manually labeled reference dataset ("3D Hyperspectral Semantic Street Scene", 3DHSSS³). The dataset has been acquired with a Riegl VZ-400i terrestrial laser scanner with co-calibrated Resonon Pika L hyperspectral line camera. The measurement setup is shown in Figure 1 (a). During scan acquisition, the laser scanner performs a 360 degree rotation, which is used to create a hyperspectral panorama image from the single lines of the hyperspectral camera. The captured points are then projected into the panorama image, using the method described above [4]. Hence, each point that is visible from the hyperspectral camera, is associated with the corresponding spectral histogram consisting of 150 spectral values between 600 nm and 1000 nm. Since the horizontal field of view of the hyperspectral camera (48.5°) is less than that of the laser scanner (100°) , not all 3D points are annotated with hyperspectral values. Areas without hyperspectral values were filtered out to make the evaluation of the results with and without hyperspectral data more comparable. As the reflected spectra are affected by external light conditions, the dataset includes scans taken at different periods in time and under varied weather conditions. This was done to enhance the stability and accuracy of the prediction models by providing training examples recorded under diverse illumination situations. The 3DHSSS dataset contains of two road sections in the city of Osnabrück, as shown in Fig. 1 (b). The first road section is located at the ICO Innovation Centre of the City of Osnabrück, Germany. The second is located near the Westerberg campus of Osnabrück University. The position difference between single scans varies from 15 to 50 meters in each scan section. The complete dataset consists of 15 different scans with an average of 45 million points each, resulting in a total of approximately one billion points. In addition to the 3D coordinates, the laser scanner returns a distance-independent reflectance measure and pulse shape deviation values for each point, which are also included.

Within these hyperspectral 3D point clouds, we labeled instances of objects from the before mentioned categories using a customized 3D labeling tool. To obtain a ground truth suitable for training, the point clouds were manually segmented into 17 different classes: Street, Sidewalk, Building, Vehicle, Street Lamp, Fence, Sign, Sign Pole, Bollard, Driveway, Low Vegetation, Bush, Tree Crown, Tree Trunk, Barrier, Distribution Box and Scan Artifact. All classes except Fence and Distribution Box are present in all scans. These two exceptions are only visible in two of the 15 scans. Furthermore, the class Tree has been divided into the subclasses Tree Crown and Tree Trunk. The reason for this is that tree crowns are often undesirable in laser scans. Especially, when surface reconstruction is done, leaves and thin structures like trunks can lead to undesirable scattered artifacts. With this purpose in mind, we introduced these separate classes to support filtering.

³ Available in the Robotic 3D Scan Repository: http://kos.informatik. uni-osnabrueck.de/3Dscans/



(a) Sensor Setup

(b) Dataset

Figure 1: Visualization sensor setup (a) and the two scenes (b) in the 3HSSS data: the hyperspectral camera (Pika L) is mounted on the terrestrial laser scanner (Riegl VZ-400i). The acquisition of the laser scan and the hyperspectral image is synchronized. Figure (b) visualizes the ICO (top) and Campus (bottom) dataset. The left column of (b) shows aerial images of the respective environments. The right column of shows the combined point clouds combined from several scan positions, colorized with the point intensities provided by the laser scanner.

4 Semantic Classification with RandLA-Net

4.1 Pre-Filtering

To prepare the hyperspectral data for classification, we perform several preprocessing steps. Since the laser scans contain many noisy points, we first apply the Statistical Outlier Removal algorithm (SOR)[32] to the recorded data to filter out artifacts. The algorithm was parameterized with the neighborhood size set to 16 and the standard deviation multiplier to 1. This setting provides a good performance in terms of outlier filtering while still preserving distant points. However, not all noise points can be eliminated with this heuristic approach. Another criterion to detect scan artifacts is to analyze the recorded intensity, as most recorded noise points have low intensity values. The problem is that some relevant materials like asphalt also absorb a lot of signal energy, hence also feature low intensity. Therefore, we decided not to apply a simple threshold filter. Instead, we examined whether these can also be reliably detected with the classification pipeline as separate class. This would allow to reliably remove them after classification without influencing the results of the relevant classes.

In a second step, the individual spectra are smoothed using the Savitzky-Golay filter. The window size was set to 5 and the degree of the polynomial to

be fitted to 3. Next, the spectra are normalized using the SNV transformation, where each spectrum is centered around zero and divided by its standard deviation. Finally, the dimension of the hyperspectral data is reduced to 64 using Principal Component Analysis.

4.2 RandLA-Net Implementation and Parameterization

In our experiments, we used the RandLA-Net implementation of Open3D-ML [33]. To integrate the hyperspectral values into the existing software, we implemented a custom dataset class in the framework, which provides functionality for loading point clouds and hyperspectral images. The 3DHSSS dataset is stored in the well known HDF5 file format in a schema similar to [34]. In addition to the 3D point clouds, hyperspectral images and label sets, all required meta information, i.e., camera parameters, resolution and scanning position, are also stored in the HDF5 file. A custom dataset class loads the corresponding training and test data from the file. Before training, the point clouds are downsampled with a voxel grid filter to ensure uniform point density (voxel size: 50 mm). After downsampling, a kd-tree is built for each point cloud to support k-nearest neighbor search. This is required to speed up the computation, since RandLA-Net analyzes the neighborhood of a point when computing features. We configured RandLA-Net to use five encoder and decoder layers. After each encoder layer, the point cloud is reduced using random sampling. The first four layers reduce the points by a factor of 4 and the last layer by a factor of 2. The output feature dimension of each layer was set to $16 \rightarrow 64 \rightarrow 128 \rightarrow 256 \rightarrow 512$. RandLA-Net receives 2^{16} points with matching hyperspectral (reduced to a dimension of 64, using PCA) and reflectance values as input. The initial reduction of the features to 16 by the first layer bears the risk that potentially important information is lost. However, with higher dimensions, the complexity of the neural network increases enormously, which makes training on consumer graphics cards unfeasible. To obtain more training data and make the model more robust, we also apply a data augmentation step before training, which adds noise and rotates or mirrors the data randomly.

We trained to RandLA-Net models on the generated training datasets: one with and one without hyperspectral data to provide a baseline for comparison. The point cloud without hyperspectral values consists of the 3D coordinates and the corresponding distance-independent reflectance values. The point clouds of the validation set are then segmented with the trained classifier and the standard metrics Accuracy, Precision, Recall, (mean) Intersection over Union are computed to assess the classification result.

5 Evaluation

In this Section we discuss the results of the trained classifiers on point clouds containing both hyperspectral and regular point clouds. Figure 2 shows exemplary classification results. Both classifiers perform well, however the hyperspectral



(a) RGB

(b) Ground truth



(c) Without hyperspectral

(d) with hyperspectral

Figure 2: Segmentation results of one point cloud. The majority of the data is well classified both with and without hyperspectral data. However, it is noticeable that without hyperspectral data (c) errors occur when classifiers similar classes such as tree (brown) and streetlamp (light blue). These can be reduced by the additional hyperspectral data (d).

information can enhance the segmentation of geometrically similar objects. One example is the classification of tree trunks which are sometimes confused with lamps without hyperspectral information. Due to the fact that vegetation can be separated easily in the spectral domain, the additional information here helps the network to learn a classifier that exploits this information as well.

For quantitative evaluation, we first compare the average values of the evaluation metrics (cf. Table 1). Because the number of samples of the various classes in the dataset is partially imbalanced, the averages were calculated using the macro method as well as the micro method. The macro method calculates the metrics for each class individually and then averages them, with equal weighting for each class. In the micro method, the individual true positives (TP), true negatives (TN), false positives (FP), and false negatives (FN) of the classes are determined and then summed up to calculate the micro statistics.

Tab. 1 clearly shows that the model with additional hyperspectral data performed better on all metrics. In addition to these averages, the classification results for each separate class are shown in Table 2. Some classes listed at the bottom of the table, such as Bollard, Sign Pole, Barrier, and Fence, are classified very poorly in some cases because the number of samples in the dataset are too low. The proportion of samples of these classes in the dataset is less than 1%. When examining the individual class results, it is evident that the use of additional hyperspectral data can improve the quality of the classifier. For the classes Scan Artifact, Street and Vehicle the hyperspectral data results in less false positives (noticeable by the higher precision values).

Scan artifacts are detected successfully with and without hyperspectral data, though without hyperspectral data the false positive rate is very high. The models can therefore also serve as a filter to remove undesired scan artifacts. As discussed before, the metrics also reflect the observation that classification of tree trunks is improved. Without hyperspectral data, Tree Trunks are often misclassified as Street Lamps (see Figure 2). Here again the hyperspectral data leads to a better distinction between similar objects with a different material composition.

6 Conclusion and Outlook

We have established an initial baseline to demonstrate that hyperspectral data as additional domain to spatial data can improve segmentation and classification with neural networks on 3D data. We showed that RandLA-Net can be extended with little effort to support the segmentation of hyperspectral 3D point clouds. The segmentation results on the provided manually labeled reference dataset prove that the inclusion of this additional data domain improves the segmentation results as expected. In future work, the data basis in terms of reference classes should be further improved to achieve a better generalization of the deep learning models. In particular, point clouds of regions containing more of the less common classes should be included. Additionally, similar classes should be combined, e.g., bush and tree crowns or barrier and bollard. Furthermore, it would also be beneficial to capture images at different seasons and more varying lighting conditions, since the spectral signature of materials is strongly dependent on the ambient illumination. Currently, we only use the standard version of RandLA-Net to include the hyperspectral data. Since hyperspectral data is high dimensional, the original network structure should be enhanced to support such data more efficiently, as it is currently designed to consider only few point at-

Table 1: Comparison of classification results for the model with and without hyperspectral data, using standard metrics: IoU (Intersection over Union),F1-Score, Accuracy, Precision and Recall. Due to imbalanced data, the averages of the metrics were calculated using both the micro and macro methods.

Spectral	Acc.	IoU (Macro)	IoU (Micro)	F1 (Macro)	F1 (Micro)	Precision (Macro)	Precision (Micro)	Recall (Macro)	Recall (Micro)
No Yes	$0.83 \\ 0.89$	$\begin{array}{c} 0.39 \\ 0.52 \end{array}$	$0.72 \\ 0.80$	$\begin{array}{c} 0.46 \\ 0.60 \end{array}$	$0.83 \\ 0.89$	$\begin{array}{c} 0.45 \\ 0.66 \end{array}$	$0.83 \\ 0.89$	$0.54 \\ 0.61$	$0.83 \\ 0.89$

		Without		With			
	Hy	perspect	ral	Hyperspectral			
Class	IoU	Pre.	Rec.	IoU	Pre.	Rec.	
Scan Artifact	0.15	0.16	0.99	0.86	0.89	0.97	
Street	0.82	0.83	0.98	0.89	0.90	0.98	
Vehicle	0.32	0.32	0.96	0.85	0.86	0.98	
Sign	0.82	0.94	0.86	0.74	0.91	0.79	
Building	0.63	0.98	0.64	0.94	0.99	0.95	
Sidewalk	0.59	0.76	0.72	0.59	0.87	0.64	
Driveway	0.03	0.08	0.04	0.12	0.17	0.29	
Tree Trunk	0.17	0.25	0.34	0.55	0.71	0.70	
Tree Crown	0.90	0.90	0.99	0.90	0.90	0.99	
Bush	0.85	0.97	0.87	0.75	0.97	0.76	
Low Vegetation	0.79	0.84	0.93	0.56	0.57	0.97	
Street Lamp	0.43	0.49	0.79	0.70	0.78	0.88	
Bollard	0.00	0.00	0.00	0.01	0.02	0.01	
Sign Pole	0.00	0.00	0.00	0.01	1.00	0.01	
Distribution Box	0.01	0.04	0.03	0.13	0.34	0.17	
Barrier	0.00	0.00	0.00	0.00	0.00	0.00	
Fence	0.01	0.02	0.04	0.15	0.25	0.27	

Table 2: Per class results of the classification models. The results are presented as Intersection over Union (IoU), Precision (Pre) and Recall (Rec)

tributes. Here, special encoder/decoder designs to pre-classify the hyperspectral data seems to be a promising approach for future developments.

References

- Guo, Y., Wang, H., Hu, Q., Liu, H., Liu, L., Bennamoun, M.: Deep learning for 3d point clouds: A survey. IEEE Transactions on Pattern Analysis and Machine Intelligence 43(12) (2021) 4338–4364
- Hu, Q., Yang, B., Xie, L., Rosa, S., Guo, Y., Wang, Z., Trigoni, N., Markham, A.: Randla-net: Efficient semantic segmentation of large-scale point clouds. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). (June 2020)
- Hackel, T., Savinov, N., Ladicky, L., Wegner, J.D., Schindler, K., Pollefeys, M.: Semantic3d. net: A new large-scale point cloud classification benchmark. arXiv preprint arXiv:1704.03847 (2017)
- Igelbrink, F., Wiemann, T., Püetz, S., Hertzberg, J.: Markerless Ad-hoc Calibration of a Hyperspectral Camera and a 3D Laser Scanner. Advances in Intelligent Systems and Computing. Springer International Publishing (2018)
- Ghamisi, P., Plaza, J., Chen, Y., Li, J., Plaza, A.J.: Advanced spectral classifiers for hyperspectral images: A review. IEEE Geoscience and Remote Sensing Magazine 5(1) (2017) 8–32
- Audebert, N., Le Saux, B., Lefèvre, S.: Deep learning for classification of hyperspectral data: A comparative review. IEEE geoscience and remote sensing magazine 7(2) (2019) 159–173

- Dargan, S., Kumar, M., Ayyagari, M.R., Kumar, G.: A survey of deep learning and its applications: A new paradigm to machine learning. Archives of Computational Methods in Engineering 27(4) (June 2019) 1071–1092
- Qi, C.R., Su, H., Mo, K., Guibas, L.J.: Pointnet: Deep learning on point sets for 3d classification and segmentation. In: Proceedings of the IEEE conference on computer vision and pattern recognition. (2017) 652–660
- 9. Qi, C.R., Yi, L., Su, H., Guibas, L.J.: Pointnet++: Deep hierarchical feature learning on point sets in a metric space. arXiv preprint arXiv:1706.02413 (2017)
- Armeni, I., Sax, A., Zamir, A.R., Savarese, S.: Joint 2D-3D-Semantic Data for Indoor Scene Understanding. ArXiv e-prints (February 2017)
- Roynard, X., Deschaud, J.E., Goulette, F.: Paris-lille-3d: A large and high-quality ground-truth urban point cloud dataset for automatic segmentation and classification. The International Journal of Robotics Research 37(6) (2018) 545–557
- Behley, J., Garbade, M., Milioto, A., Quenzel, J., Behnke, S., Stachniss, C., Gall, J.: Semantickitti: A dataset for semantic scene understanding of lidar sequences. In: Proc. of the IEEE/CVF International Conf. on Computer Vision (ICCV). (2019)
- Manolakis, D., Lockwood, R., Cooley, T.: Hyperspectral Imaging Remote Sensing. Cambridge University Press (2016)
- 14. Xue, J., Su, B.: Significant remote sensing vegetation indices: A review of developments and applications. Journal of sensors **2017** (2017)
- Wu, D., Sun, D.W.: Advanced applications of hyperspectral imaging technology for food quality and safety analysis and assessment: A review-part i: Fundamentals. Innovative Food Science & Emerging Technologies 19 (2013) 1–14
- Van der Meer, F.D., Van der Werff, H.M., Van Ruitenbeek, F.J., Hecker, C.A., Bakker, W.H., Noomen, M.F., Van Der Meijde, M., Carranza, E.J.M., De Smeth, J.B., Woldai, T.: Multi-and hyperspectral geologic remote sensing: A review. International Journal of Applied Earth Observation and Geoinformation 14(1) (2012) 112–128
- 17. Goetz, A.F.: Three decades of hyperspectral remote sensing of the earth: A personal view. Remote Sensing of Environment **113** (2009) S5–S16
- Hege, E.K., O'Connell, D., Johnson, W., Basty, S., Dereniak, E.L.: Hyperspectral imaging for astronomy and space surveillance. In: Imaging Spectrometry IX. Volume 5159., International Society for Optics and Photonics (2004) 380–391
- Vaiphasa, C.: Consideration of smoothing techniques for hyperspectral remote sensing. ISPRS Journal of Photogrammetry and Remote Sensing 60(2) (2006) 91–99
- Raczkowska, M.K., Koziol, P., Urbaniak-Wasik, S., Paluszkiewicz, C., Kwiatek, W.M., Wrobel, T.P.: Influence of denoising on classification results in the context of hyperspectral data: High definition ft-ir imaging. Analytica Chimica Acta 1085 (2019) 39–47
- Savitzky, A., Golay, M.J.E.: Smoothing and Differentiation of Data by Simplified Least Squares Procedures. Anal. Chem. 36(8) (July 1964) 1627–1639
- 22. Green, A., Berman, M., Switzer, P., Craig, M.: A transformation for ordering multispectral data in terms of image quality with implications for noise removal. IEEE Transactions on Geoscience and Remote Sensing 26(1) (1988) 65–74
- Lee, J., Woodyatt, A., Berman, M.: Enhancement of high spectral resolution remote-sensing data by a noise-adjusted principal components transform. IEEE Transactions on Geoscience and Remote Sensing 28(3) (1990) 295–304
- Rinnan, A., van den Berg, F., Engelsen, S.B.: Review of the most common preprocessing techniques for near-infrared spectra. TrAC Trends in Analytical Chemistry 28(10) (2009) 1201–1222

- Barnes, R.J., Dhanoa, M.S., Lister, S.J.: Standard normal variate transformation and de-trending of near-infrared diffuse reflectance spectra. Applied Spectroscopy 43(5) (1989) 772–777
- Geladi, P., MacDougall, D., Martens, H.: Linearization and scatter-correction for near-infrared reflectance spectra of meat. Appl. Spectrosc. 39(3) (May 1985) 491– 500
- Buckley, S.J., Kurz, T.H., Howell, J.A., Schneider, D.: Terrestrial lidar and hyperspectral data fusion products for geological outcrop analysis. Computers & Geosciences 54 (2013) 249–258
- Weinmann, M.: Fusion of hyperspectral, multispectral, color and 3d point cloud information for the semantic interpretation of urban environments. International Archives of the Photogrammetry, Remote Sensing & Spatial Information Sciences (2019)
- Brell, M., Segl, K., Guanter, L., Bookhagen, B.: 3d hyperspectral point cloud generation: Fusing airborne laser scanning and hyperspectral imaging sensors for improved object-based information extraction. ISPRS Journal of Photogrammetry and Remote Sensing 149 (2019) 200–214
- 30. Roy Choudhury, M., Das, S., Christopher, J., Apan, A., Chapman, S., Menzies, N.W., Dang, Y.P.: Improving biomass and grain yield prediction of wheat geno-types on sodic soil using integrated high-resolution multispectral, hyperspectral, 3d point cloud, and machine learning techniques. Remote Sensing 13(17) (2021)
- 31. Goodfellow, I., Bengio, Y., Courville, A.: Deep Learning. MIT Press (2016) http: //www.deeplearningbook.org.
- Rusu, R.B., Marton, Z.C., Blodow, N., Dolha, M., Beetz, M.: Towards 3d point cloud based object maps for household environments. Robotics and Autonomous Systems 56(11) (2008) 927–941
- Zhou, Q., Park, J., Koltun, V.: Open3d: A modern library for 3d data processing. CoRR abs/1801.09847 (2018)
- Wiemann, T., Igelbrink, F., Pütz, S., Hertzberg, J.: A file structure and reference data set for high resolution hyperspectral 3d point clouds. IFAC-PapersOnLine 52(8) (2019) 403–408

12