Roof Surface Classification with Hyperspectral and Laser Scanning Data – An Assessment of Spectral Angle Mapper and Support Vector Machines

Stephanie BRAND

Abstract

The urban environment is characterised by a variety of different surface materials. For the discrimination of urban materials, hyperspectral imaging proved a valuable tool. In this study, two methods for classification, Spectral Angle Mapper and Support Vector Machines, are compared on a hyperspectral dataset to derive a detailed map of roof materials. Spectral similarity of different materials, especially with low reflectance and no distinct absorption features can complicate the classification process. Therefore, hyperspectral data were supplemented with laser scanning data to not only discriminate roof from ground data but also to use roof inclination to distinguish roof spectra. A binary roof mask from a laser scanning dataset was used to restrict the classification to roofs only. After testing the two classifiers on this reduced dataset, the approach was extended by incorporating inclination information in the classification process. Comparison between the classified images is done visually and quantitatively using confusion matrices. It can be shown that both classifiers are suitable for the classification of roof materials with the Spectral Angle Mapper results yielding higher classification accuracies than Support Vector Machines. For both classification approaches, the confusion between several materials was reduced by the incorporation of roof inclination, thus improving overall accuracy.

1 Introduction

Hyperspectral imaging comprises the measurement and analysis of reflectance spectra collected in small, contiguous spectral bands. A range of studies with various applications showed the usefulness of hyperspectral data for classification purposes, e.g. urban material mapping. In this study, two different pixel-based classifiers, Spectral Angle Mapper (SAM) and Support Vector Machines (SVM) are compared to assess their performance in the mapping of roof surfaces. In the urban environment, where large variations in illumination occur, SAM should be well suited because it has proven to be insensitive to these effects. SAM has been used in an urban context, for example to map roof materials on their vulnerability for hailstorms (BHASKARAN et al 2001). SVM has been used in land cover mapping with hyperspectral and multispectral data and gave better classification results than other classifiers (HUANG et al 2002, PAL & MATHER 2004). Supplementing hyperspectral data with elevation information to enhance classification by reducing spectral confusion between urban land cover types has been successfully employed by MADHOK & LANDGREBE (1999).

Jekel, T., Car, A., Strobl, J. & Griesebner, G. (Eds.) (2012): GI_Forum 2012: Geovizualisation, Society and Learning. © Herbert Wichmann Verlag, VDE VERLAG GMBH, Berlin/Offenbach. ISBN 978-3-87907-521-8.

A building mask derived from the DSM of the laser scanning dataset was used in this study to limit the classification process to roofs only. LEMP & WEIDNER (2005) included inclination information of roofs in a segment-based classification process of roof surfaces. They proved that the supplemented information led to an improvement of accuracy. Roof inclination should therefore be useful in the classification process of roof surfaces that show similar reflectance curves and are therefore hard to distinguish. The inclination information could help in the classification process if the confused materials occur on differently inclined roofs. As a consequence, the study focuses on the following points: Which classification method gives better results in the mapping of roof surfaces? Can the incorporation of roof inclination improve the accuracy of the classification results?

2 Study Area and Data

The study area is the main campus of the KIT (Karlsruhe Institute of Technology) in the city of Karlsruhe, Germany. The area is approximately 1km × 0.6km in size containing a mixture of roof materials of various ages and conditions. The buildings of the campus area are framed by residential buildings in block development (see figure 1). The main materials are red roofing tile and slate on residential-like houses, gravel and stone slab on flat roofs, as well as various metals on industrial-like buildings. Advantageous to the study is the fact that the buildings on the campus are usually larger than residential buildings; this means that the larger the roofs, the more pure pixels on a roof can be expected.

Hyperspectral data from the HyMap sensor were acquired from the Institute of Photogrammetry and Remote Sensing (IPF) in July 2003 during the HyEurope campaign organised by the DLR (German Aerospace Center). The HyMap sensor is an airborne spectrometer which consists of 128 bands in a nearly contiguous wavelength spectrum from 0.44µm to 2.5µm with a nominal bandwidth of 15-20 nm. The spatial resolution of the dataset is 4m × 4m. Pre-processing was undertaken by the DLR, including atmospheric correction with ATCOR4, conversion to apparent reflectance and geocorrection of the dataset, using the DSM from the laser scanning dataset provided by the IPF. During pre-processing, two bands were excluded from the dataset resulting in 126 bands.

The laser scanning data were acquired in March 2002 with the TopoSys II system. The generation of the DSM was done using first pulse and last pulse data and converted to 1m × 1m pixels. A binary building mask was then derived to discriminate buildings and non-buildings (see figure 2). A second dataset with the inclination of roofs in percent was also calculated from the DSM.

Aerial imagery from spring 2001 was used to gain knowledge about the roof surface materials in the study area and to derive correct and representative training areas. For accuracy assessment, a vector reference dataset of the campus roof materials was created, displayed in figure 3. The roof materials of about 60 building complexes were determined by aerial imagery and field checks. Additionally, reflectance curves of roof materials in the image were compared to roof spectra described in literature.

Fig. 1: Hyperspectral image of study area with main campus (framed)

Fig. 2: Binary roof mask derived from DSM (right); Buildings displayed in white

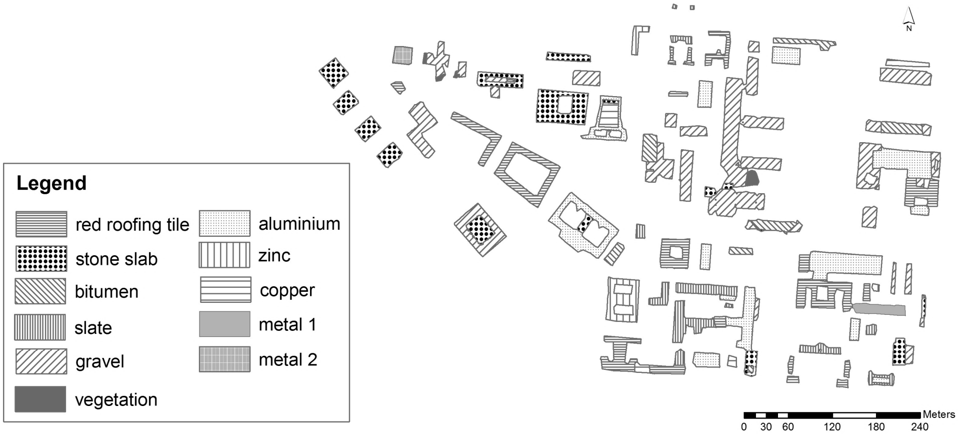

Fig. 3: Vector reference dataset with roof surface materials

3 Methodology

In the following, the methods used in this study are described, from data preparation to the classification approaches with SAM and SVM.

3.1 Data preparation

The hyperspectral dataset was first investigated for high level of noise in the image bands using the homogeneous area method (SMITH & CURRAN 1996). If the noise in the signal response is high, the image band will not reliably represent the feature of interest. The investigation resulted in the removal of bands 1, 63-66 and 95, leaving 120 bands in the dataset.

After clipping the hyperspectral dataset to the size of the study area, the image was resampled to 1m × 1m resolution using nearest-neighbour interpolation. Thus it was ensured that the hyperspectral dataset had the same pixel size as the roof mask. Afterwards, the roof mask from laser scanning data was applied to the hyperspectral dataset to mask out all areas that are not building roofs. With this dataset the first classification processes with SAM and SVM were undertaken.

As described above, one task of this study was to analyse whether the incorporation of slope information improves the classification results. Two masks of flat roofs and inclined roofs from the available slope dataset from the DSM were created. Different definitions exist in the literature to determine which angle sets the threshold between flat and inclined roofs. As the main task of the roof mask is to distinguish between inclined roofs covered with e.g. roofing tiles or slate and flat roofs covered with gravel, stone slab or bitumen, the definition of the minimum angle for inclined roofs was taken which states that the lowest achievable incline with tile roofs is 11° (WORMUTH et al. 2007). Therefore, the threshold between flat and inclined roofs was set to 11° or approximately 18%. Figure 4 displays the two final roof inclination masks. Following this, the hyperspectral dataset was clipped with the slope masks to provide two datasets for classification where buildings of the opposite inclination are excluded. These were used in the classification process with slope information.

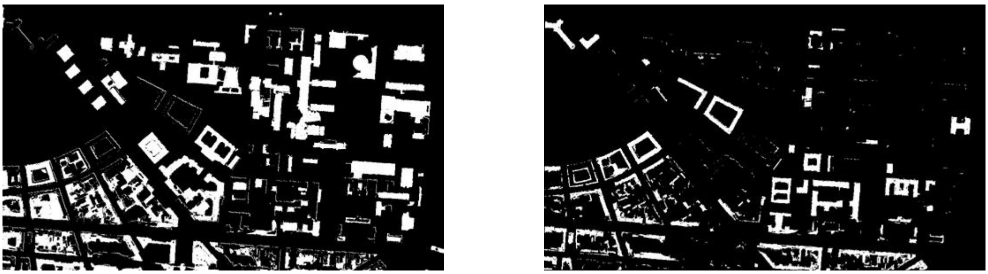

Fig. 4: Final masks of flat roofs (left) and inclined roofs (right). Roofs displayed in white.

Roof material	Characteristics	Geometry	
		flat	inclined
Red tile	Distinct absorption features at 0.52µm, 0.67µm and 0.87µm		+
Slate	Low reflectance; no distinct absorption features; spectrum similar to bitumen		+
Stone slab	Spectrum similar to gravel; few absorption features	+	
Gravel	Spectrum similar to stone slab; few absorption features	+	
Bitumen	Low reflectance; spectrum similar to slate	+	
Vegetation	Reflectance peak at 0.55µm; steep ascent of reflectance to the NIR at 0.63µm	+	+
Aluminium	High reflectance; characteristic absorption band at 0.84µm	+	+
Zinc	Low reflectance; distinct absorption feature at 1.02µm and 0.52µm	+	+
Copper	Reflectance peak at 0.55µm; broad absorption band around 0.73µm	+	+
Metal 1	Medium reflectance; distinct absorption feature at 0.57µm	+	+
Metal 2	Medium reflectance; Several absorption bands, strongest at 0.84µm	+	+

Fig. 5: Overview of roof materials with description of characteristics and geometry

As both SAM and SVM are supervised classification methods where the classifier needs to be trained, representative training data of the roof materials in the image were derived. For the subsequent classification, the same training datasets were used to ensure comparability of the results. Training classes were selected for every roof material. Some materials required the definition of more classes to fully represent the material in the image due to differences in age, colour, weathering and illumination.

The training classes were extracted from parts of the roofs that were as homogeneous as possible, not including chimneys, roof windows or other materials. However, because of the spatial resolution of the hyperspectral data of 4m and the nature of the roofs, the training classes contained a mixture of relatively pure and mixed pixels. Altogether 20 classes were derived which represent 11 roof materials (see figure 5). Two different metals (metal 1 and 2) were included for which no name of the material could be determined but which are spectrally unique and were thus added to the training areas.

3.2 Classification with SAM

The Spectral Angle Mapper (SAM) uses spectral similarity to allocate pixel in classes. The spectral similarity is determined by calculating the spectral angle between the reference spectra from the training classes and each image spectrum. The spectra are treated as vectors in n-dimensional space where n is given by the number of bands. The image spectrum is assigned to the training class spectrum to which it has the smallest angle. The advantage of SAM is that is considered insensitive to illumination and albedo effects (KRUSE et al. 1993) that occur due to shade on an object or the facing of the roof towards or away from the recording sensor, factors that have high influence in urban areas. Dark and bright illuminated pixels are treated equally, with darker pixels being situated nearer the origin than brightly illuminated pixels. The angle distance to the vector of the training class spectrum will stay the same, meaning that a pixel of the same material under different illumination conditions will most likely be classified in the same class.

The Spectral Angle Mapper of the software ENVI 4.7 was used for classification. A maximum angle threshold in radians is used as input defining the maximum acceptable angle between a training and a pixel vector. Any pixel with an angle larger than the specified threshold is not classified. After testing different angle thresholds for the image and also different angle thresholds for each class in the image, one angle threshold was found to classify the image best. Using different angles for each class, this did not improve the accuracy but resulted in overclassification of some roof materials. A small angle of 0.1 rad left too many pixels unclassified (about 25%), 0.3 rad showed underrepresentation of some roof materials and 0.5 rad provided the best results and was used in the classification.

3.3 Classification with SVM

Support Vector Machines are derived from the field of machine learning theory and have already been applied successfully in several studies (HUANG et al. 2002, PAL & MATHER 2004). The classifier is well suited for high-dimensional data as used in this study due to the fact that it does not assume certain statistical class distributions. The basic concept of SVM classification is to fit an optimal hyperplane between classes using training samples at the edge of the class distributions, the so called support vectors. The simplest case is a two class problem with linearly separable classes in an n-dimensional space. Theoretically, many hyperplanes could be fitted to separate the two classes, but there is only one optimal plane where all pixels of a class are located on the same side of the plane and the distance from each class to the hyperplane is largest. If the classes are not linearly separable, a slack variable is introduced which indicates the distance of the pixel from the "correct" side of the plane. A penalty term C is controlling the magnitude of the penalty for pixels that lie on the wrong side of the plane. It is tried to maximise the margin between the classes while penalising the pixels on the wrong side of the plane using the parameter C. If classes are not linearly separable, the dataset is mapped to a higher dimensional space using a kernel function. Thus, a separating linear plane can be fitted to divide the classes. A variety of different kernels can be used for this task. One kernel which is widely applied, also in the software package used in this study, is the Radial Basis Function (RBF kernel). The parameter γ needs to be specified to control the width of the Gaussian kernel. How well SVM performs is dependent on the magnitude of the parameters C and γ. The larger either of the parameters, the higher is the risk of over-fitting to the training data, thus providing a poor generalisation.

For the Support Vector Machine approach the software imageSVM was used (VAN DER LINDEN et al. 2009). ImageSVM uses the Gaussian radial basis function kernel (RBF kernel). After scaling of the image to a range between 0 and 1 (necessary to provide suitable values for the parameterisation of the classifier), the parameters γ and C are searched from a range of values. The pairs of γ and C are tested against each other to find the parameters with the best performance using a grid search (cross-validation). With the best parameters found, the classification process is then executed.

Parameter search for γ and C for the RBF-kernel was done using a 5-fold cross validation. Several searches were undertaken using different search ranges for the grid-search. Several pairs of γ and C resulted in a high cross-validation accuracy of 99.9198%. The model with the smallest γ value and the same high cross validation accuracy was chosen. This is because the value γ controls the width of the kernel. If γ is small, each support vector has a large area of influence on other points. Thus the risk of overfitting is greater with higher γ values because only the training areas are correctly classified. This led to the final parameters of $\gamma = 2^{-13}$ and $C = 2^{20}$ which were used for classification.

3.4 Incorporation of slope

To answer the question whether the incorporation of roof inclination improves the classification result for each classifier, slope information was added to the classification process. This was done using a two-step approach: Input for the classification with each classifier were two hyperspectral "slope" datasets, one which only contained flat roofs, one which contained inclined roofs. Next, the training areas were divided into those occurring on flat roofs (stone plates, gravel, bitumen), those on inclined roofs (slate and tiles) and those that are not specific to a certain inclination and which were therefore included in both training class groups (see also figure 5 for the assignment of the materials to an inclination type). The two hyperspectral "slope" datasets with the corresponding training classes were then classified separately with the two classifiers. The classification results of each classifier were afterwards combined again to receive one final classified image with 11 classes. Figure 6 displays the workflow of the classification process with roof inclination.

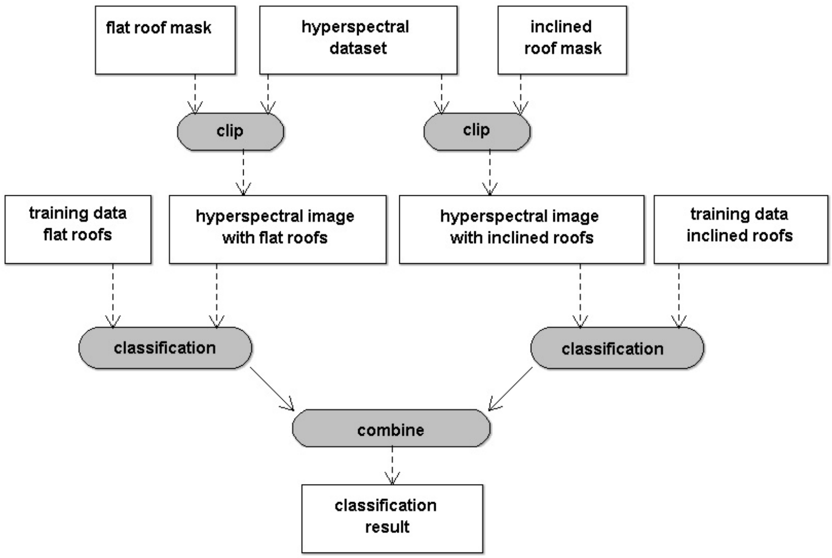

Fig. 6: Classification workflow with roof inclination

4 Results

In the following, the results are first analysed individually and then compared with each other. The analysis was done manually and quantitatively with confusion matrices using the reference dataset displayed in figure 3. To achieve a good evaluation of the accuracy of the results, around 45% of the roof pixels in the image were used in the accuracy assessment. Finally, the classification results with the slope information are described and analysed.

4.1 Classification without roof inclination

Both classifiers used in this study are pixel-based and classify each pixel according to its spectral response without regarding neighbouring pixels. Therefore, the classification results show a pixellated appearance and most roofs are not homogeneous. The overall accuracy of SAM was better with 70.48% than that of SVM (59.06%), also showing higher Kappa indices (SAM 0.64; SVM 0.52).

In general, red tile roofs were classified well in the SAM classification results, not only on large roofs like those on the campus but also on smaller residential building roofs. Gravel and stone plate, however, are often confused. The confusion is higher in the SAM than the SVM result. This can be attributed to the fact that their material composition is very similar and thus are their reflectance curves. Slate roofs also prove to be problematic as they are often confused with bitumen roofs and vice versa. Here, the SAM performs slightly better with a producer's accuracy of 59% and user's accuracy of 31%, whereas SVM yields only 37% and 13% respectively. Aluminium roofs are difficult to classify for SAM. In the SAM result, some aluminium roofs are confused with other materials, especially copper. This leads to a low producer's accuracy of 59%. This is contrary to the result of SVM, which classifies aluminium quite well with a producer's accuracy of 71%. A reason for the bad performance of the aluminium class with SAM might be that some pixels exceed the sensor capabilities and therefore the aluminium signature is not recognisable anymore for SAM. For the results of all classes refer to Table 1.

Table 1: Classification accuracy by class for SAM and SVM. Positive difference between SAM and SVM is marked in bold letters (i.e. SVM classified better)

	Producer's accuracy (in %)			User's accuracy (in %)		
Classes	**SAM**	**SVM**	**SVM-SAM**	**SAM**	**SVM**	**SVM-SAM**
Copper	77.10	71.19	-5.91	42.39	70.01	**27.62**
Tile	88.72	75.26	-13.46	95.96	96.64	**0.68**
Stone slab	56.86	54.90	-1.96	69.22	50.22	-19.00
Zinc	67.12	58.61	-8.51	97.35	89.05	-8.30
Slate	67.58	68.01	**0.43**	47.59	29.59	-18.00
Gravel	78.00	43.64	-34.36	76.41	29.59	-6.09
Bitumen	59.54	37.83	-21.71	31.59	13.30	-18.29
Metal 1	92.72	82.63	-10.09	81.70	92.75	**11.05**
Aluminium	54.63	71.48	**16.85**	89.54	91.27	**1.73**
Vegetation	60.97	19.17	-41.8	19.57	39.30	**20.33**
Metal 2	63.71	81.45	**17.74**	23.10	20.14	-2.96

As can be expected in an urban environment, the amount of shadow is high. This is also the case in the study area, especially in the area of the campus with high building complexes casting a shadow on adjacent roofs. Here the SAM approach offers better classification results as well. The shadow pixels where the spectral response is attenuated are classified more correctly than with SVM. The SVM classifier assigns these shadow pixels with low reflectance to classes with low reflectance and no distinct absorption bands like bitumen and slate.

4.2 Classification with roof inclination

Theoretically, the incorporation of roof inclination information in the classification process should improve the accuracy of the classifications because some materials only occur on inclined roofs, others only on flat roofs. This should lead to a less pronounced mix-up of these materials when the classification of flat and inclined roofs is done separately. As was shown in chapter 4.1, the materials bitumen and slate are very often confused, as well as gravel and slate. Gravel and bitumen only occur on flat roofs, slate only on inclined roofs. In the classification with the slope masks from the DSM, the accuracy was however not improved compared to the classification result without roof inclination. The producer's and user's accuracy of slate and bitumen did indeed improve, indicating that the roof inclination helps in the classification process.

This result was analysed and it was found that the lack of improvement is attributed to the fact that the slope mask was not correct for some buildings: some roofs, which are in reality inclined, have low inclination values and were thus included in the mask of flat roofs. And the opposite was true for some flat roofs. This meant for example that a tile roof, which was correctly classified as tile in the classification without inclination information, was wrongly assigned to the class stone slab or bitumen in the classification with inclination information.

The classification was therefore repeated with a manually created slope mask which corrected the inclination errors, to test if an improvement in classification accuracy would then be visible. The overall accuracies for the SAM and SVM result improved by about 4.5%, to 75% for SAM and 64% for SVM respectively. Table 2 gives an overview of the overall accuracies with and without inclination.

Table 2: Classification accuracies in comparison with and without inclination information

Classification	Overall accuracy (%)	Kappa-Coefficient
SAM	70.48	0.64
SAM slope original	70.05	0.64
SAM slope manual	74.95	0.70
SVM	59.06	0.52
SVM slope original	59.12	0.52
SVM slope manual	63.83	0.58

The roofs which were previously assigned to the wrong slope dataset were correctly classified again. For the SAM result, the improvement was greatest for slate which gained about 20% in producer’s and 26% in user’s accuracy. User’s accuracy for bitumen and stone slab was improved by about 7%. The accuracy for bitumen and slate also increased for the SVM result, gaining 36% (bitumen) and 28% (slate) in producer’s accuracy.

In general it can be said, that the classification results are less pixellated and the roofs are clearer and more homogeneously assigned to their class when inclination is used in the classification process. The problem of the classification of aluminium in the SAM approach and the problem with the shadowed roof in the SVM approach, however, remain. As these problems do not occur because of confusion between materials with similar reflectance curves on differently inclined roofs, the incorporation of roof inclination does not help to solve these problems.

5 Summary and Conclusion

In this study, two classification approaches, Spectral Angle Mapper (SAM) and Support Vector Machines (SVM), were compared to assess the ability of mapping roof materials with a hyperspectral dataset in the city of Karlsruhe. Classification was done with 11 roof material classes on a hyperspectral dataset which was clipped to include only buildings. The Spectral Angle Mapper thereby provided better results than Support Vector Machines, contrary to studies that compared SVM to other classifiers (HUANG et al. 2002, PAL & MATHER 2004). The overall accuracy was about 11% higher for SAM than for SVM. The strength of the SAM approach lay in the classification of shadowy areas where it performed better than SVM. SVM, on the other hand, proved more stable in the classification of materials with high reflectance like aluminium, whose response sometimes even exceeded sensor capabilities. Additionally, a thesis was tested which assumes that inclination information in the classification process improves the accuracy of the results. First, a dataset with flat roofs and one with inclined roofs was classified and afterwards the results combined. It was evident that the accuracy of the classification result did not improve. The reasons for this were errors in the inclination dataset: parts of inclined roofs were contained in the flat roof mask and vice versa thus diminishing classification accuracy. By applying a manual slope mask which corrected for the errors in the original masks, classification accuracy was increased, proving that the original thesis was correct. The classification accuracy of the SAM and SVM results were increased by about 4.5%, showing an increase of the accuracy of some classes like slate of more than 20% with both methods.

This study showed that both classifiers are suited to map roof surfaces. Further work in this respect could be done by testing different SVM parameters to improve these classification results. Another interesting approach would be to incorporate the slope information directly using data fusion methods and test the performance of the two classifiers.

Acknowledgments

Thanks to Dr. Uwe Weidner from the Institute of Photogrammetry and Remote Sensing, KIT, for providing the datasets and the helpful discussions.

References

BHASKARAN, S., FORSTER, B. & NEAL, T. (2001), Integrating Airborne Hyperspectral Sensor Data with GIS for Hail Storm. Post-Disaster Management. 22nd Asian Conference on Remote Sensing 5-9, November 2001, Singapore. http://www.crisp.nus.edu.sg/~acrs2001/pdf/178bhask.pdf (last accessed 2012-01-29).

HUANG, C., DAVIS, L. S. & TOWNSHEND, J. R. G. (2002), An assessment of support vector machines for land cover classification. Int. Journal of Remote Sensing, 23 (4), 725-749.

KRUSE, F. A., LEFKOFF, A. B., BOARDMAN, J. B., HEIDEBRECHT, K. B., SHAPIRO, A. T., BARLOON, P. J., & GOETZ, A. F. H. (1993), The Spectral Image Processing System (SIPS) – Interactive Visualization and Analysis of Imaging Spectrometer Data. In: Remote Sensing of Environment, Special issue on AVIRIS, 44, 145-163.

LEMP, D. & WEIDNER, U. (2005), Improvements of roof surface classification using hyperspectral and laser scanning data. Proceedings URBAN 2005, Arizona (CD-ROM).

MADHOK, V. & LANDGREBE, D. (1999), Supplementing hyperspectral data with digital elevation. IGARSS 1999, Hamburg (CD-ROM).

PAL, M. & MATHER, P. M. (2004), Assessment of the effectiveness of support vector machines for hyperspectral data. Future Generation Computer Systems, 20, 1215-1225.

SMITH, G.M. & CURRAN, P. J. (1996), The signal-to-noise ratio (SNR) required for the estimation of foliar biochemical concentrations. Int. Journal of Remote Sensing, 17 (5), 1031-1058.

VAN DER LINDEN, S., RABE, A., OKUJENI, A. & HOSTERT, P. (2009), imageSVM classification, Application Manual: imageSVM version 2.0. Berlin.

WORMUTH, R., DIERKS, K. & FLEISCHMANN, H. D. (2007), Dächer. In: DIERKS, K. & WORMUTH, R. (Eds.), Baukonstruktion. 6. Auflage. Neuwied, Werner-Verlag.