



Road condition estimation using deep learning with hyperspectral images: detection of water and snow

Daniil Valme*, Javier Galindos and Dhanushka Chamara Liyanage

Department of Electrical Power Engineering and Mechatronics, Tallinn University of Technology, Ehitajate tee 5, 19086 Tallinn, Estonia

Received 21 March 2023, accepted 16 June 2023, available online 15 February 2024

© 2024 Authors. This is an Open Access article distributed under the terms and conditions of the Creative Commons Attribution 4.0 International License CC BY 4.0 (<http://creativecommons.org/licenses/by/4.0>).

Abstract. Road surface condition monitoring is one of the most crucial tasks for vehicle perception systems. The presence of water, snow, ice, or any other substance covering the road surface directly affects the rolling resistance and controllability of the vehicle, which is directly related to the safety of the traffic participants. Many sensors, such as RGB cameras, infrared sensors, and mmWave sensors, are used to monitor and inspect road surfaces. The research aims to provide a tool to segment an input image into correct classes. The DeepHyperX toolbox was used for the rapid prototyping of deep learning (DL) classification models for hyperspectral images. The effectiveness of the developed algorithm in several case studies is presented, and it is verified that the low number of iterations is enough to detect the water, snow, and ice on the road surface.

Keywords: road surface condition monitoring, water, snow, advanced driver assistance system.

1. INTRODUCTION

In recent years, the issue of improving road safety has become increasingly relevant. According to the WHO *Global Status Report on Road Safety 2018*, the number of annual road traffic deaths has reached 1.35 million. Road infrastructure greatly influences the causes of fatal and severe injuries in traffic collisions [1]. The recent developments in the automotive industry aim to propose solutions for this problem. Current cars are outfitted with advanced driver assistance systems (ADAS) to safeguard drivers better and reduce accidents [2]. The poor performance of autonomous vehicles and driver assistance systems in inclement weather, such as snow, rain, hail, and fog, is now one of the most critical problems in developing these technologies [3]. For that reason, monitoring the road surface condition is crucial for traffic safety [4].

ADAS systems can be significantly improved by real-time road friction estimates (RFE) [5]. The performance

of these systems strongly depends on the characteristics of the sensors used in the detection system. Cameras, LiDARS, and radars are the primary sensors used in autonomous driving to gather environmental data. Modern ADAS systems implement simultaneous localization and mapping (SLAM) technology and use different types of sensors, such as RGBD cameras, LiDARS, and sonars. RGB cameras are the most popular due to their low price and the ability to provide comprehensive data [6]. The widespread use of RGB cameras in visual SLAM was driven by advancements in central processing unit (CPU) and graphics processing unit (GPU) technology. In [5], the front-view RGB camera sensor images are used for road surface condition (RSC) classification. Wet/water, slush, snow, and ice RSC classifications are determined to have discrete regions of interest through texture-based features from the drivable surface, sky, and surrounds. In other papers, the RGB image was converted into a spatial frequency spectrum, and the frequency distribution corresponding to different road surface types was subsequently analysed [5,7].

* Corresponding author, daniil.valme@taltech.ee

The hyperspectral camera is one of the most promising and rapidly evolving types of sensors. Compared to RGB cameras, hyperspectral cameras capture more information than the visible light range in the electromagnetic spectrum. Those cameras can capture visible light, near infrared (NIR), short-wave infrared (SWIR), and medium-wave infrared (MWIR), making cameras less sensitive to changing environmental conditions. The 3D hyperspectral image (HSI), also referred to as a hypercube, contains hundreds of contiguous spectral bands. The reflectance characteristics at different wavelength bands produce a unique spectral signature. Based on this signature, image pixels of similar materials can be identified [8]. Further post-processing of the hypercube enables the prediction of the objects' materials. Information about materials may be essential for making decisions in the context of autonomous driving. In HSI classification tasks, two main challenges exist: the large spatial variability of spectral signatures and the limited available training samples versus the high dimensionality of hyperspectral data. The first challenge is often caused by many factors, such as changes in illumination and environmental, atmospheric, and temporal conditions. The second challenge results in ill-posed problems for some methods and reduces the generalization ability of classifiers [9].

Despite the challenges mentioned above, the topic of applying HSI for vehicles is drawing attention. For example, in [10], the authors presented a manually annotated dataset containing the HSI of a structured urban terrain. Another HSI dataset presented in [11] includes road images acquired in urban and rural scenes. Authors in [8] focused on implementing an HSI for pictures captured in unstructured off-road terrain with grass, trees, and water puddles. The goal of the proposed HSI-based method for pixel-wise image annotation was to reduce the workload of human annotation.

The goal of this research was to develop a framework that allowed for the quick testing of different deep learning (DL) models on HSIs. The next step was to segment an input image into the correct classes as a result of the inference of the proposed models. This paper's main focus was on water and snow detection, as this information can potentially provide the ADAS systems with the next layer of information.

2. MATERIALS AND METHODS

2.1. Data collection

The images used in the datasets were captured using the Specium IQ mobile hyperspectral camera. The camera captures the HSIs in the 400–1000 nm range with a resolution of 512×512 pixels using the push-broom

image acquisition method with a line scan imaging sensor. Specium IQ uses a CMOS sensor that provides information regarding the 204 spectral bands. Specium IQ lens has a focal length of 21 mm, a field of view (FOV) 31×31 degrees, and a minimum object distance for the lens of 150 mm [12].

The HSI dataset was collected in the surrounding metropolitan area of Tallinn University of Technology in Estonia. The images with snow were taken in February, while the rest were taken in September. Data was primarily collected when it was cloudy outside and out of direct sunshine. The dataset contains 35 water and 26 snow HSIs with no defined ground truth. The HSIs come in a .hdr and .dat format, which both need to be imported to obtain the information from the files.

2.2. Data exploration

The studied dataset includes RGB images in different resolutions (1280×960 px and 512×512 px), as well as a raw HSI datafile in .dat (512×512 px). An HSI taken with the Specim IQ mobile hyperspectral camera contains 204 bands. Therefore, 204 images were taken at different wavelengths. In comparison, the RGB representation of an image only contains three colour channels. Thus, using hyper-spectral data enables access to more data about a particular scene. The spectrum for different semantic classes is presented in Appendix B. In particular, the spectrum for asphalt (Appendix B1), concrete (Appendix B2), and water (Appendix B3) is shown in the corresponding graphs. It is noteworthy to observe the variance of the spectrum within the same semantic class. In Fig. 1, the average spectrum for each semantic class is presented. One should observe the similarity between classes like asphalt and water to understand the challenge of distinguishing between these two classes.

2.3. Data processing and band reduction

The structure of the adopted framework, DeepHyperX [13], only allows to analyse one image at a time, performing pixel-wise classification using a default patch of surrounding pixels to assist in classification. Usually, HSI datasets, such as Pavia University [14] or Indian Pines [15], only consist of one image. However, in the current research, multiple images were made (35 images containing water and 26 snow). Therefore, all HSI arrays were concatenated into one big 'image' to make the DeepHyperX framework applicable to this problem. This method works correctly as every pixel is classified by itself regardless of the surrounding image.

The output shape of the concatenated matrix is $512 \times 17920 \times 204$, containing more than 1.8 billion 32-bit float values inside a NumPy array that requires approximately




Fig. 1. The average spectrum of different classes.

12 GB of memory. To conform with the DeepHyperX formatting, the training data (i.e. hyperspectral data) should be in a .mat format, and the ground-truth labels should be in a .tiff format. A .mat file can store a limited number of values, which is more than three times less capacity than is required for the water dataset. Due to the high number of bands contained in an HSI, the data files are often computationally unworkable. Thus, a dimensionality reduction technique is needed to deal with these large data files due to computational restrictions. According to the work presented in [16], feature extraction (transforming/combining different bands) is preferred over feature selection (elimination of image bands). This work also states that principal component analysis (PCA), non-negative matrix factorization (NMF), and locally linear embedding (LLE) are the optimal techniques to increase the overall accuracy of the model – 30 out of 200 components are sufficient for PCA/NMF, and 90 bands for LLE. The disadvantage of PCA is that more than 99% of the variability is explained in the first three components, which means the other 27 components used explain less than 1% of the variability.

In the current research work, PCA was used to reduce dimensionality from 204 to 30 bands. Working with all the bands is computationally very challenging since the data exceeds 12 GB, being beyond the limits of most commercial RAMs. After performing PCA, the data is significantly (more than six times) smaller and computationally more efficient.

2.4. Data labelling

Data labelling is the process of identifying raw data and adding one or more meaningful and informative labels to provide context so that the machine learning model

can learn from it. The labelling was performed using the Labelbox training data platform [17]. This platform was selected due to its built-in collaboration and management features, powerful toolkit for object annotation, and user-friendly interface.

The dataset was composed of scenes captured in the urban environment, including roads and parking lots in different conditions. For labelling, there were 14 classes defined, each defining a specific type of object present in a particular scene represented with the corresponding colour (see Table 1). For labelling time reduction, the following simplifications were made:

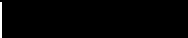













- asphalt road in all conditions was labelled as ‘Asphalt’;
- urban environment objects such as sculptures, street-lights, containers, road signs, bicycles, parking spots, etc., were labelled as ‘Object’;
- all the buildings were labelled as ‘Concrete’ despite the presence of other materials in their structure;
- all the combinations of ice, mud, and snow in the same area were considered as ‘Ice’;
- trees and bushes were labelled as ‘Trees’.

The most accurate labelling was done for the areas of the images where there was water, ice, or snow present because the quality of labelling of those classes is crucial for achieving reliable results. It should be mentioned that classes such as ‘Trees’ or ‘Sky’ were given less attention than the ones mentioned above as the research was mainly focused on the road surface.

2.4.1. Water dataset labelling

The main feature of this dataset is the variety of puddles combined with asphalt roads in different conditions. This allows the investigation of the hyperspectral signatures of different wet asphalt surfaces. According to the char-

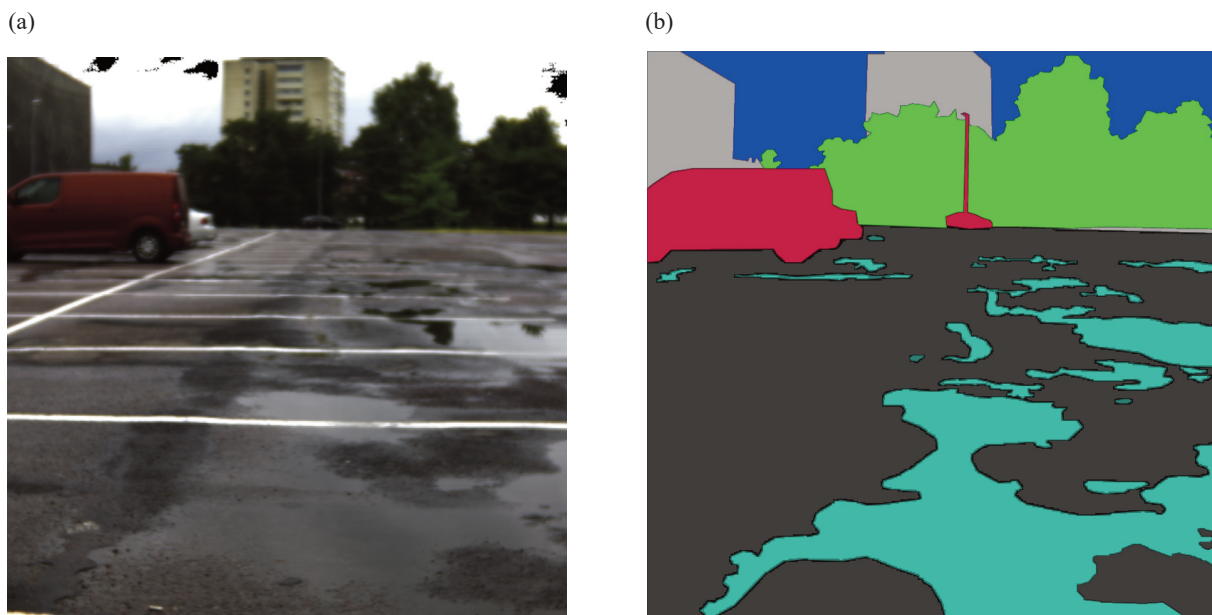
Table 1. Labelling classes, label colours, and HEX code values

No.	Class	Red	Green	Blue	HEX code	Colour sample
1.	Undefined	0	0	0	#000000	
2.	Grass	0	102	0	#006600	
3.	Concrete	170	170	170	#aaaaaa	
4.	Asphalt	64	64	64	#404040	
5.	Trees	0	255	0	#00ff00	
6.	Rocks	110	22	138	#6e168a	
7.	Water	68	187	170	#44bbaa	
8.	Sky	0	0	255	#0000ff	
9.	Gravel	187	134	51	#bb8833	
10.	Object	192	6	64	#c00640	
11.	Dirt	230	230	30	#e6e61e	
12.	Mud	99	66	34	#634222	
13.	Snow	240	240	240	#f0f0f0	
14.	Ice	120	170	230	#78aac6	

acteristics of the road and its condition, the asphalt class can typically be expressed as a combination of subclasses: old asphalt, new asphalt, road seams filled with tar, asphalt covered with paint, and potholes repaired with stone crumbs. All areas with a significant presence of water on the road surface were considered as water. The example of labelling the scene, including the road covered with water, is shown in Fig. 2.

2.4.2. Snow dataset labelling

In the snow dataset, there were multiple areas of the road in which the snow was combined with asphalt, water, ice, and mud. Thus, it was essential to make the annotations as precise as possible. In the scenes presented in the dataset, the white areas of the images mainly corresponded to snow, and all other areas were considered as ice (including

**Fig. 2.** Example of performed water labelling: (a) source image and (b) labelled image.

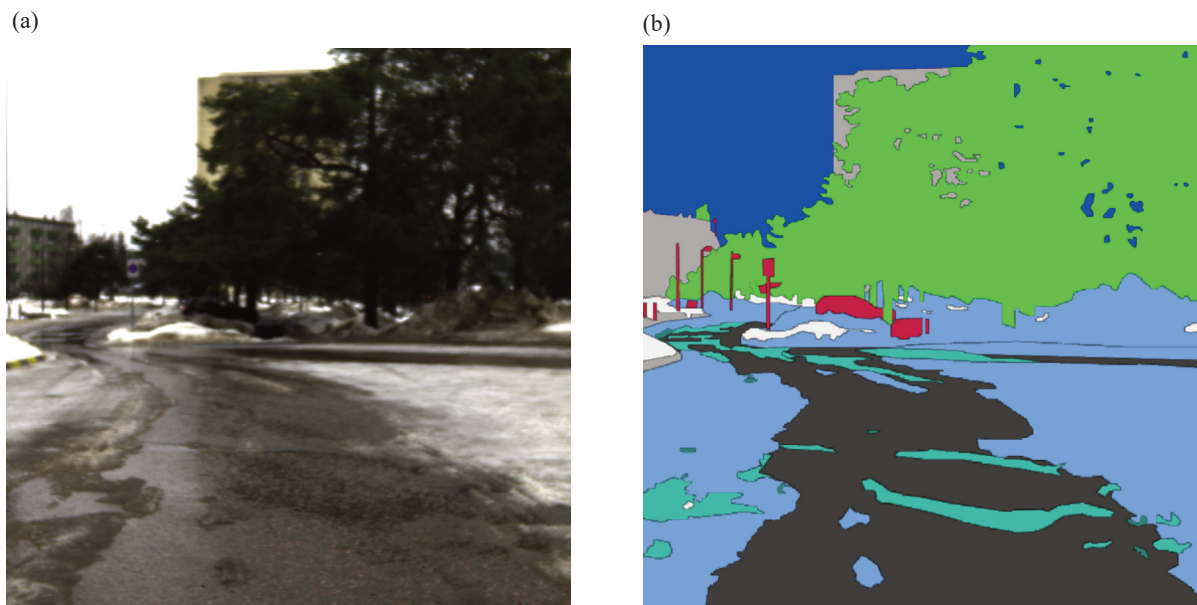


Fig. 3. Example of performed snow labelling: (a) source image and (b) labelled image.

the combination of ice, snow, and mud). Fig. 3 illustrates an example of labelling the scene that includes snow.

2.5. Model validation

Due to the high number of bands in a hyperspectral image, the data files are often computationally unworkable. Thus, a dimensionality reduction technique is needed to deal with these large data files due to computational restrictions.

Despite the class-balancing modification in the loss function, the segmentation tasks are usually an unbalanced problem. As the number of pixels per a semantic class is significantly different, the accuracy could be biased toward estimating the model's goodness. Thus, the F1-score is preferred to validate the model's goodness, which is a balanced combination of precision and recall. The class imbalance for the water dataset is well represented by the examples shown in Fig. 4. As shown in the results, a model without class balancing leads to a higher overall accuracy (72.7%) than the model with class balancing (46.7%). However, this information could be misleading regarding the goodness of the model; in particular, the F1-score of water for the first model (without class balancing) is 19%, whereas the F1-score of the class-balanced model is 37%. Similar phenomena are observed for the minority semantic classes.

3. RESULTS

All hyperspectral images and ground-truth labels were concatenated into a single image, resulting in a $512 \times (\text{amount of images}) \times 30$ shape for the training data and a $512 \times (\text{number of images})$ shape for the ground-truth labels. The results are obtained through pixel-wise classification using a default patch size of 5×5 surrounding pixels to assist in classification. The 30 channels result from performing PCA on the 204 hyperspectral bands. Using the DeepHyperX toolbox, earlier mentioned models were trained in parallel to observe the model response. Due to hardware constraints, not all available images were used in the training and validation process. The input images were split into 85% for training data and 15% for validation data. The specifications of the hardware used for the processing are described in Appendix A.

The validation data are pixels from the training image which have been left out and are unknown to the model during training. Class balancing was performed on the dataset due to the issues mentioned in Section 2. Not performing class balancing leads to overfitting, showing an acceptable overall accuracy but a very poor F1-score for the critical classes (see Table 2).




Fig. 4. Imbalance in data points per class: example with a little water within the scene (a), example with no water within the scene (b).

Table 2. F1-scores of water dataset without data balancing, accuracy 72.7%

No.	Class	F1
1.	Grass	0.092
2.	Concrete	0.673
3.	Asphalt	0.838
4.	Trees	0.652
5.	Water	0.198
6.	Sky	0.534
7.	Object	0.234
8.	Mud	0.410

Table 3. F1-scores of water dataset with balancing, accuracy 46.7%

No.	Class	F1
1.	Grass	0.103
2.	Concrete	0.494
3.	Asphalt	0.547
4.	Trees	0.516
5.	Water	0.373
6.	Sky	0.440
7.	Object	0.313
8.	Mud	0.252

3.1. Water

The models were run for ten epochs each for visualization, and the intermediate results were compared. Using 30 hyperspectral images and their corresponding ground truth, the 3D convolutional neural network (CNN) model developed by Hamida et al. [18] yields the most accurate results within the scope of the performed experiments with an overall validation accuracy of 46.7% and an F1-score for the water of 37.3% (see Table 3). The loss and validation accuracy plots of the water dataset are shown in Fig. 5a,b, respectively.

The corresponding confusion matrix in Fig. 6 shows the (expected) confusion between asphalt and water. Based on the results displayed in the confusion matrix, it is possible to conclude that the asphalt class is most often

incorrectly recognized. The class of asphalt is most frequently confused with water, which is very reasonable. HSI is a surface inspection method, so wet or water-covered surfaces can be recognized as water. Since some of the images in the dataset were taken immediately after the rain, in some cases the water does not cover the surface in a continuous layer, as is the case with large puddles, but fills in small depressions, forming micro-puddles. In addition to water, classes of grass, concrete, and trees are erroneously found on the asphalt surface. These assumptions can be considered partially true since the labelling was done in a very general way.

Observing the predictions versus the ground truth in Fig. 7, the approximate area where water is present is semi-accurately predicted. In general, it can be said that the approximate location of certain classes can be de-




Fig. 5. (a) Water dataset training process loss graph, (b) water dataset validation process accuracy graph.




Fig. 6. Confusion matrix of the water dataset.

terminated; however, it is not possible to talk about recognizing the shapes of objects. In addition, one can see significant errors in the prediction regarding the example of the ‘Sky’ class, which can also be found on the road. It is important to note that this result was achieved at 10 epochs. Increasing the number of epochs leads to a better result; an example with 50 epochs is presented in the next section. However, for real-life applications, knowing the approximate area is more critical than pixel-accurate predictions of water.

3.2. Snow

Compared to the HSIs including water, in the case of snow, the models were run for 50 epochs each, and the




Fig. 7. 3D CNN model developed by Hamida et al. [18] applied for 10 epochs, representing ground truth (a1–a5) versus prediction (b1–b5) for five sample water HSIs. (Continued on the next page)









Fig. 7. Continued on the next page

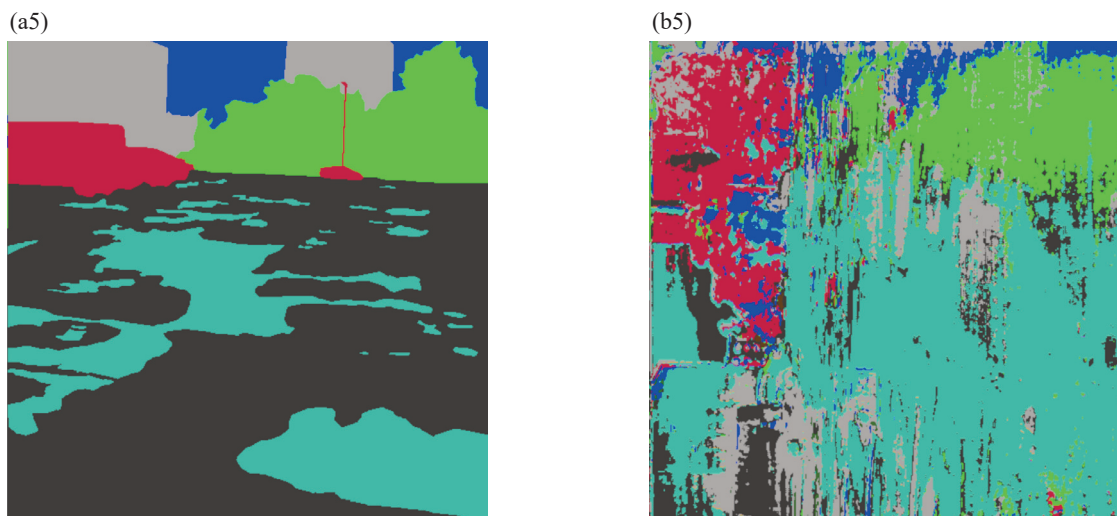


Fig. 7. Continued

Table 4. F1-scores of the snow dataset, accuracy 65.9%

No.	Class	F1
1.	Grass	0.082
2.	Concrete	0.833
3.	Asphalt	0.848
4.	Trees	0.672
5.	Water	0.516
6.	Sky	0.538
7.	Object	0.735
8.	Mud	0.092
9.	Snow	0.485
8.	Ice	0.848

intermediate results were compared to each other. Using ten hyperspectral images and their corresponding ground truth, the 3D CNN model developed by Hamida et al. [18] yields the most accurate results within the scope of the performed experiments with an overall validation accuracy of 65.9% and an F1-score for snow of 48.5% (see Table 4). Another useful outcome is the F1-score for the ice with a value of 84.8%. Extraction of this information may also be useful from the perspective of the ADAS systems and driving safety.

The loss and validation accuracy plots of the snow dataset are shown in Fig. 8a,b, respectively. The corresponding confusion matrix in Fig. 9 shows the (expected) confusion between mud and snow, and the model prediction versus the ground truth is observed in Fig. 10. The confusion matrix shows that the snow class often gets confused with the mud class. The possible reason for that




Fig. 8. (a) Snow dataset training process loss graph, (b) snow dataset validation process accuracy graph.




Fig. 9. Confusion matrix of the snow dataset.

may be the fact that during the labelling process, all types of snow were defined in a general group including pure loose snow, snow mixed with mud, sand, etc. For that reason, the dirty snow could be recognized as snow. In the future, this aspect can be further developed by introducing subclasses for the main snow class.

It is important to note that the original images also included a large area where the ice class was represented. Snow drifts covered with a solid crust of ice were also noted as ice. Due to the differentiation of these classes during the labelling process, surfaces covered with ice

were very well predicted. For example, in Fig. 10b1, it can be seen that the model was able to successfully distinguish between water, asphalt, and ice in the same zone. The output images of the model run for 50 epochs can already provide more precise information regarding both the shapes of the objects and their location.

4. DISCUSSION

The hardware limitations are a clear bottleneck in HSI classification problems. The size of each HSI is around 200 MB. The state-of-the-art frameworks, such as DeepHyperX, allow for training with only one image. Therefore, there was a need to concatenate the images to have a more extensive training sample, which is computationally very challenging on most high-end consumer hardware. It is noteworthy that the training process requires specific advanced hardware. In the current study, high-performance hardware through remote machines was used. The training duration of the models on high-end GPUs took more than 24 hours. However, better results could be obtained using area-wise classification instead of pixel-wise classification. Using only one pixel in the classification process is prone to error due to distortions in capturing the scene through a hyper-spectral camera. Using patches for classification instead of only one pixel is beneficial in this application, given that water usually exists in a collection of multiple pixels (pools of water) rather than as a single observed pixel at a time. Although Hamida et al. [18] in the DeepHyperX toolbox use a default patch size of 5 and do not observe




Fig. 10. 3D CNN model developed by Hamida et al. [18] applied for 50 epochs, representing ground truth (a1–a5) versus prediction (b1–b5) for five sample snow HSIs. (Continued on the next page)




Fig. 10. Continued on the next page

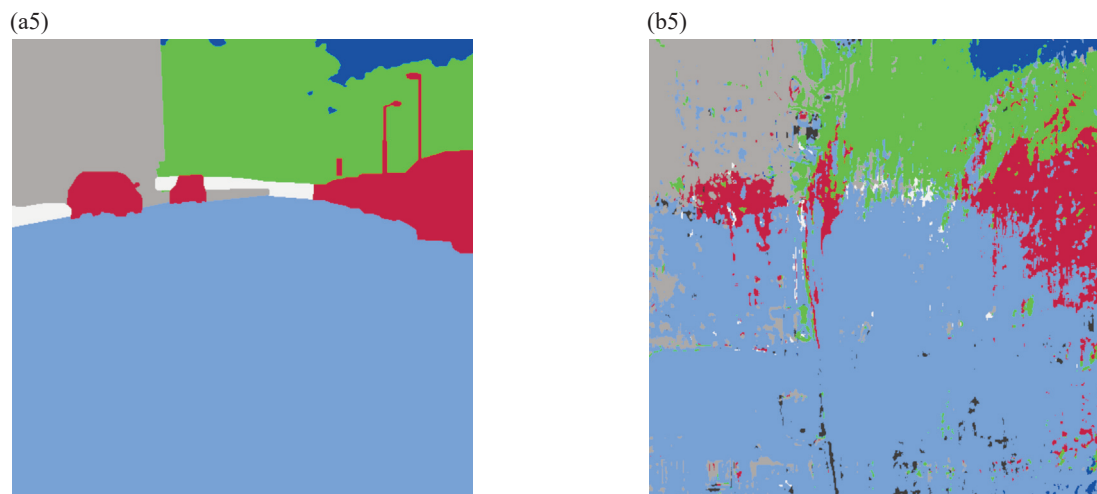


Fig. 10. Continued

only one pixel at a time, the patch size was not explored in further detail. Future research should investigate the optimal patch size for this classification problem.

Furthermore, there is a lack of online documentation regarding HSI classification/segmentation compared to other topics of DL. The absence of proven models demonstrates that techniques such as transfer learning or fine-tuning cannot be easily applied. These techniques are beneficial in problems like the one this work faces, where the training sample is not big enough to train a deep neural network from scratch and have good results. Further research should consider the application of transfer learning for this problem.

Also, to achieve better results in future tests, more detailed image labelling should be considered. This can include creating additional classes to extend the existing ones; for example, the snow class can be split into several types of snow. In this way, it will be possible to obtain more detailed information regarding the materials of the captured objects. In addition to this, it is worth considering pixel-wise labelling as a tool to increase accuracy. To improve recognition accuracy, it may be necessary to expand the dataset with more diverse images.

Finally, further improvements should be considered regarding the loading of data instead of concatenating all images. Concatenating images while using pixel patches for classification leads to a lot of noise at the boundary/cross of every image with the next image. It could influence the effectiveness of the classification model.

5. CONCLUSIONS

The DeepHyperX toolbox was used for the rapid prototyping of DL hyperspectral classification models. To comply with the pixel-wise classification demands of the toolbox,

every HSI and the corresponding ground truth should be concatenated into a single image each. Therefore, 30 images of the water dataset were concatenated, resulting in a single image with a shape of $512 \times 15360 \times 204$. Similarly, ten images of the snow dataset were concatenated, resulting in a single image with a shape of $512 \times 5120 \times 204$. The corresponding ground-truth segmented images underwent the same concatenation process for the water and snow datasets.

Both datasets were separately reduced in dimension by performing PCA, resulting in 30 channels instead of 204. The 3D CNN by Hamida et al. [18] is the best model provided by the DeepHyperX toolbox, resulting in an accuracy of 46.7% and an F1-score of 37.3% for the water dataset run with ten epochs, and an accuracy of 65.9% and an F1-score of 48.5% for the snow dataset run with 50 epochs.

Due to the imbalanced class distribution of the images, class balancing was performed to avoid overfitting. Model predictions show that the model confuses asphalt and water. However, comparing the predictions to the base truth, the approximate areas where water was present were semi-accurately predicted. For real-life applications, knowing the approximate area is more critical than pixel-accurate predictions of water. As a result, it can be concluded that water, snow, and ice on the road can be detected using HSI classification.

ACKNOWLEDGEMENTS


The research was supported by the Estonian Research Council under the grant PSG453 ‘Digital twin for propulsion drive of autonomous electric vehicle’. The publication costs of this article were covered by the Estonian Academy of Sciences.

APPENDIX A


Hardware specifications

Task	CPU	GPU	RAM
Pre-processing	Intel Core i7 9th Generation 9750H (2.60 GHz)	Nvidia GeForce RTX 2060 6 GB GDDR6	16 GB
Training process	AMD Threadripper 3970X 32-Core/64-Thread Processor (3.7 GHz)	Nvidia 3090 with 24 GB of graphics memory	128 GB


APPENDIX B



B1. Spectrum analysis of asphalt class.



B2. Spectrum analysis of concrete class.



B3. Spectrum analysis of water class.

REFERENCES

1. WHO. *Global Status Report on Road Safety 2018*. <https://www.who.int/publications/i/item/9789241565684> (accessed 2024-01-24).
2. Ziebinski, A., Cupek, R., Grzechca, D. and Chruszczyk, L. Review of advanced driver assistance systems (ADAS). *AIP Conf. Proc.*, 2017, **1906**(1), 120002. <https://doi.org/10.1063/1.5012394>
3. Zang, S., Ding, M., Smith, D., Tyler, P., Rakotoarivelo, T. and Kaafar, M. A. The impact of adverse weather conditions on autonomous vehicles: how rain, snow, fog, and hail affect the performance of a self-driving car. *IEEE Veh. Technol. Mag.*, 2019, **14**(2), 103–111. <https://doi.org/10.1109/MVT.2019.2892497>
4. Munawar, H. S. Image and video processing for defect detection in key infrastructure. In *Machine Vision Inspection Systems: Image Processing, Concepts, Methodologies and Applications* (Malarvel, M., Nayak, S. R., Panda, S. N., Pattnaik, P. K. and Muangnak, N., eds). Wiley Online Library, 2020, 159–177. <https://doi.org/10.1002/9781119682042.CH7>
5. Roychowdhury, S., Zhao, M., Wallin, A., Ohlsson, N. and Jonasson, M. Machine learning models for road surface and friction estimation using front-camera images. In *Proceedings of the 2018 International Joint Conference on Neural Networks (IJCNN), Rio de Janeiro, Brazil, 8–13 July 2018*. IEEE, 2018, 1–8. <https://doi.org/10.1109/IJCNN.2018.8489188>
6. Chen, C., Zhu, H., Li, M. and You, S. A review of visual-inertial simultaneous localization and mapping from filtering-based and optimization-based perspectives. *Robotics*, 2018, **7**(3), 45. <https://doi.org/10.3390/ROBOTICS7030045>
7. Shinmoto, Y., Takagi, J., Egawa, K., Murata, Y. and Takeuchi, M. Road surface recognition sensor using an optical spatial filter. In *Proceedings of Conference on Intelligent Transportation Systems, Boston, MA, USA, 12 November 1997*. IEEE, 1997, 1000–1004. <https://doi.org/10.1109/ITSC.1997.660610>
8. Liyanage, D. C., Hudjakov, R. and Tamre, M. Hyperspectral imaging methods improve RGB image semantic segmentation of unstructured terrains. In *Proceedings of the 15th International Conference Mechatronic Systems and Materials (MSM), Bialystok, Poland, 1–3 July 2020*. IEEE, 2020, 1–5. <https://doi.org/10.1109/MSM49833.2020.9201738>
9. Li, S., Song, W., Fang, L., Chen, Y., Ghamisi, P. and Benediktsson, J. A. Deep learning for hyperspectral image classification: an overview. *IEEE Trans. Geosci. Remote Sens.*, 2019, **57**(9), 6690–6709. <https://doi.org/10.1109/TGRS.2019.2907932>
10. Basterretxea, K., Martínez, V., Echanobe, J., Gutiérrez-Zaballa, J. and Del Campo, I. HSI-drive: a dataset for the research of hyperspectral image processing applied to autonomous driving systems. In *Proceedings of the 2021 IEEE Intelligent Vehicles Symposium (IV), Nagoya, Japan, 11–17 July 2021*. IEEE, 2021, 866–873. <https://doi.org/10.1109/IV48863.2021.9575298>
11. Lu, J., Liu, H., Yao, Y., Tao, S., Tang, Z. and Lu, J. Hsi road: a hyper spectral image dataset for road segmentation. In *Proceedings of the 2020 IEEE International Conference on Multimedia and Expo (ICME), London, UK, 6–10 July 2020*. IEEE, 2020, 1–6. <https://doi.org/10.1109/ICME46284.2020.9102890>
12. Specim. Specim IQ Technical Specifications. <https://www.specim.fi/iq/tech-specs/> (accessed 2022-12-14).
13. Audebert, N., Saux, B. and Lefèvre, S. Deep learning for classification of hyperspectral data: a comparative review. *IEEE Geosci. Remote Sens. Mag.*, 2019, **7**(2), 159–173. <https://doi.org/10.1109/MGRS.2019.2912563>
14. Fauvel, M., Chanussot, J., Benediktsson, J. A. and Sveinsson, J. R. Spectral and spatial classification of hyperspectral data using SVMs and morphological profiles. *IEEE Trans. Geosci. Remote Sens.*, 2008, **46**(11P2), 3804–3814.

15. Li, J., Bioucas-Dias, J. M. and Plaza, A. Semisupervised hyperspectral image segmentation using multinomial logistic regression with active learning. *IEEE Trans. Geosci. Remote Sens.*, 2010, **48**(11), 4085–4098. <https://doi.org/10.1109/TGRS.2010.2060550>
16. Rychlewski, D. *Hyperspectral image classification of satellite images using compressed neural networks*. Master's thesis, 2020. https://www.researchgate.net/publication/344817503_Hyperspectral_Image_Classification_of_Satellite_Images_Using_Compressed_Neural_Networks (accessed 2022-12-05).
17. Labelbox. <https://labelbox.com/> (accessed 2022-12-05).
18. Ben Hamida, A., Benoit, A., Lambert, P. and Ben Amar, C. 3-D deep learning approach for remote sensing image classification. *IEEE Trans. Geosci. Remote Sens.*, 2018, **56**(8), 4420–4434. <https://doi.org/10.1109/TGRS.2018.2818945>

Teeolude hindamine sügavõppe abil hüperspektraalsete piltidega: vee ja lume tuvastamine

Daniil Valme, Javier Galindos ja Dhanushka Chamara Liyanage

Teeolude jälgimine on sõidukite tajusüsteemide üks olulisemaid ülesandeid. Vee, lume, jää või teekatet katva muu aine olemasolu mõjutab sõiduki veeretakistust ja juhitavust, mis on otseselt seotud liikluses osalejate ohutusega. Andureid, nagu RGB-kaamerad, infrapunaandurid ja millimeeterlaineline radarid, kasutatakse teekatete jälgimiseks ja kontrollimiseks. Uurimuse eesmärk on pakkuda tööriista sisendpildi segmenteerimiseks õigetes klassidesse. Tööriista DeepHyperX kasutati sügavõppe klassifitseerimismudelite kiireks prototüüpimiseks hüperspektraalsete piltide jaoks. Esitletakse arendatud algoritmi tõhusust mitmesugustes juhtumianalüüsid ja leiab kinnitust, et väike iteratsioonide arv on piisav vee, lume ja jää tuvastamiseks teekatte pinnal.