

# On-chip hyperspectral image segmentation with fully convolutional networks for scene understanding in autonomous driving<sup>☆</sup>

Jon Gutiérrez-Zaballa<sup>a,\*</sup>, Koldo Basterretxea<sup>a</sup>, Javier Echanobe<sup>b</sup>, M. Victoria Martínez<sup>b</sup>, Unai Martínez-Corral<sup>a</sup>, Óscar Mata-Carballeira<sup>a</sup>, Inés del Campo<sup>b</sup>

<sup>a</sup> Department of Electronics Technology, University of the Basque Country, 48013, Bilbao, Spain

<sup>b</sup> Department of Electricity and Electronics, University of the Basque Country, 48940 Leioa, Spain

## ARTICLE INFO

Dataset link: <https://ipaccess.ehu.es/HSI-Drive/>, <https://gitlab.com/EHU-GDED/FLOSS/HSI-Drive>

### Keywords:

Hyperspectral imaging  
scene understanding  
fully convolutional networks  
autonomous driving systems  
system on chip  
benchmarks

## ABSTRACT

Most of current computer vision-based advanced driver assistance systems (ADAS) perform detection and tracking of objects quite successfully under regular conditions. However, under adverse weather and changing lighting conditions, and in complex situations with many overlapping objects, these systems are not completely reliable. The spectral reflectance of the different objects in a driving scene beyond the visible spectrum can offer additional information to increase the reliability of these systems, especially under challenging driving conditions. Furthermore, this information may be significant enough to develop vision systems that allow for a better understanding and interpretation of the whole driving scene. In this work we explore the use of snapshot, video-rate hyperspectral imaging (HSI) cameras in ADAS on the assumption that the near infrared (NIR) spectral reflectance of different materials can help to better segment the objects in real driving scenarios. To do this, we have used the HSI-Drive 1.1 dataset to perform various experiments on spectral classification algorithms. However, the information retrieval of hyperspectral recordings in natural outdoor scenarios is challenging, mainly because of deficient color constancy and other inherent shortcomings of current snapshot HSI technology, which poses some limitations to the development of pure spectral classifiers. In consequence, in this work we analyze to what extent the spatial features codified by standard, tiny fully convolutional network (FCN) models can improve the performance of HSI segmentation systems for ADAS applications. In order to be realistic from an engineering viewpoint, this research is focused on the development of a feasible HSI segmentation system for ADAS, which implies considering implementation constraints and latency specifications throughout the algorithmic development process. For this reason, it is of particular importance to include the study of the raw image preprocessing stage into the data processing pipeline. Accordingly, this paper describes the development and deployment of a complete machine learning-based HSI segmentation system for ADAS, including the characterization of its performance on different embedded computing platforms, including a single board computer, an embedded GPU SoC and a programmable system on chip (PSoC) with embedded FPGA. We verify the superiority of the FPGA-PSoC over the GPU-SoC in terms of energy consumption and, particularly, processing latency, and demonstrate that it is feasible to achieve segmentation speeds within the range of ADAS industry specifications using standard development tools.

## 1. Introduction

Today, thanks to the availability of small-size, portable, snapshot hyperspectral cameras, it is possible to set up HSI processing systems on moving platforms. The use of drones for precision agriculture and ecosystem monitoring is probably one of the most active and mature application domains of this technology [1]. The research into how hyperspectral information can be used to develop more capable and

robust ADAS is, on the contrary, in its infancy. HSI provides rich information about how materials reflect the light of different wavelengths (spectral reflection), and this can be used to identify and classify surfaces and objects in a scene [2,3]. In fact, RGB images tend to suffer from metamerism (two objects with different reflection spectra but having the same color under a certain light source), so HSI can help to resolve this phenomenon and become a powerful solution for

<sup>☆</sup> This work was partially supported by the Basque Government under grants KK-2021/00111 and PRE\_2021\_1\_0113 and by the Spanish Ministry of Science and Innovation under grant PID2020-115375RB-I00. We thank the University of the Basque Country for its allocation of computational resources.

\* Corresponding author.

E-mail address: [j.gutierrez@ehu.es](mailto:j.gutierrez@ehu.es) (J. Gutiérrez-Zaballa).

<https://doi.org/10.1016/j.sysarc.2023.102878>

Received 23 October 2022; Received in revised form 7 April 2023; Accepted 13 April 2023

Available online 28 April 2023

1383-7621/© 2023 The Author(s). Published by Elsevier B.V. This is an open access article under the CC BY-NC-ND license (<http://creativecommons.org/licenses/by-nc-nd/4.0/>).

object segmentation [4]. Thus, with the application of appropriate information processing techniques, HSI can help to enhance the accuracy and robustness of current ADAS for object identification and tracking and, eventually, can be used for scene understanding, which is a step forward in the achievement of more capable and intelligent ADAS/ADS (Autonomous Driving Systems) [5,6].

HSI segmentation of real driving scenes is, however, very challenging for a variety of reasons that are relative to both the lack of control over the environment (lighting, weather, distance of objects etc.) and the technology of portable snapshot HSI cameras, i.e. the use of spatial mosaic and tiled spectral filters, the limited spectral range of CMOS sensors and the lack of a photometer-based exposure-time setting. First, acquiring HSI outdoors implies working with varying illumination and environmental conditions and dealing with the presence of moving objects and changing backgrounds, which strongly conditions the setup of the acquisition system. In consequence, acquired images need to undergo a time-consuming preprocessing stage to transform a 2D image with radiance data into a 3D cube of reflectance values which, under these variable conditions and changing camera setups, can hardly guarantee the required robustness for the subsequent information processing. Secondly, the spectral signatures of the different objects of interest for ADAS/ADS in a driving scene may show a weak spectral separability when acquired under the conditions described herein. Thirdly, the extraction of additional spatial features that could help to improve the segmentation of items with similar spectral reflectance is problematic as a result of the enormous diversity of shapes, view angles, distances and scales in an image. Finally, ADS/ADAS development implies executing image processing and AI algorithms onboard and in real-time (the latter meaning with a sufficiently low latency), and this must be performed on resource constrained hardware platforms with tight power consumption restrictions. This entails not only evaluating the performance indexes (recall/precision) of the inferred classification or segmentation maps but carefully analyzing the computational complexity of the algorithms involved and how to efficiently implement them on the target devices.

In this work, which is an extended version of [7], we describe an ML approach based on FCN for the segmentation of HSI video recorded in real driving conditions as part of a research project which aims at exploring the applicability of small-size snapshot hyperspectral sensor technology to the enhancement of vision-based ADAS. This is a multidisciplinary work that includes investigating in an integrated manner HSI processing techniques, machine learning and AI algorithms and models, and processor and advanced SoC design techniques to develop a functional real-time image segmentation system that can realistically be implemented in an embedded processing system. We present two application examples of the design of a HSI segmentation system that could be applied to ADS/ADAS. These examples are used throughout the paper to show the system development process, from the algorithmic design, (i.e. ML model training and optimization), to the final system prototyping and testing on different embedded processing systems: embedded multicore microprocessors, an embedded GPU SoC, and a PSoC with an embedded FPGA. The first, simpler example aims at the segmentation of images by the pixelwise classification of three categories: tarmac, road marks and “rest” or non-drivable areas, which could be used to enhance automatic lane keeping and trajectory planning systems for ADS. The second example incorporates two additional categories, vegetation and sky, to make a total of five different categories to be classified, which extends the system’s ability for a broader scene understanding. The aim was to evaluate, first, to what extent deep convolutional segmentation models such as U-Net, which perform successfully in image segmentation tasks in other application fields, can help to overcome the limitations revealed in previous stages of our research of purely spectral classifiers. The second aim was to evaluate whether it is pertinent to try out the implementation of such models into embedded processing platforms, considering the real-time requirements of the application.

To do so, we have trained and evaluated the segmentation algorithms under study using a HSI database specifically created for this purpose [8]. Various FCN variants have been used in order to analyze their performance on the experimental data. However, in this approach we have avoided complex, compound FCN models to keep the model sizes under control. Moreover, the selected FCN model hyperparameters have finally been optimized in order to achieve a good trade-off between segmentation accuracy and model complexity to reduce computing time (FLOPS) and memory footprint (reduced number of parameters). As stated before, small-size hyperspectral snapshot cameras require quite a complex preprocessing pipeline to get the 3D cubes containing spectral reflectance data from the recorded radiance raw images. For an ADS/ADAS application, this preprocessing must be carried out on-the-fly and on the onboard processing device. In addition, we include an ablation study to observe the effect of those techniques separately on each of the stages of the preprocessing. Finally, we have conducted a comprehensive benchmarking to exhaustively characterize the prototype performances in terms of latency, power and energy consumption, both at the image preprocessing stage and at the inference of the deployed network. Besides this, the benchmarking also considers the comparison between the execution of the FCN models using only the processing systems (microprocessors) with the combined execution of the processing system and the dedicated coprocessors for AI inference (GPU and FPGA), which comprises a previous quantization of the models. In fact, the effect of the quantization process is also addressed to check how it affects the previously mentioned measurements as well as the memory footprint and the segmentation accuracy

The main contributions of this paper are summarized as follows:

- The spectral data analysis and the development of HSI segmentation algorithms for scene understanding in ADAS are performed on images obtained with a snapshot camera in real driving scenarios and under diverse lighting and weather conditions. This guarantees that the results obtained are realistic and can be extrapolated to an eventual real implementation based on this same technology.
- Some inherent limitations of the spectral information obtained from multispectral filter-array based HSI snapshot cameras are identified and addressed in the development of HSI segmentation algorithms.
- A lightweight FCN is developed for the efficient segmentation of HSI images for scene understanding in ADAS. This improves the performance of pure spectral classifiers while remaining simple enough to be deployed successfully on an embedded computing platform.
- Real HSI segmentation prototypes are implemented and their performance characterized in terms of memory footprint, latency, power and energy consumption by the deployment of the whole processing pipeline on different commercial embedded heterogeneous computing platforms: an embedded multicore microprocessor, an embedded GPU SoC, and a PSoC with embedded FPGA. This involves analyzing not only the FCN inference stage for image segmentation but also the raw image preprocessing stage that precedes the formation of the hyperspectral cubes.

The rest of the paper is organized as follows: Section 2 contains the most relevant works regarding Deep Learning-based segmentation with either RGB or HSI in ADAS as well as articles where the superiority of HSI over RGB is shown. In Section 3 we present the experimentation setup, including an analysis of the database which shows the implications of acquiring HSI images outdoors and in challenging environments. Some initial results with baseline spectral are included too, where a comparison between segmentation with HSI or pseudoRGB is included. Section 4 includes the FCN model development for which a detailed hyperparameter tuning process has been carried out. Besides this, it also covers the way in which the training has been conducted and presents the segmentation results. Section 5 contains the

deploying of both the FCN and the cube preprocessing algorithm on the three embedded computing platforms whose performance is deeply benchmarked in terms of latency, power and energy consumption, resource usage and accuracy. Section 6 concludes the paper and discusses possible future work.

## 2. Related work

There is very limited prior work on ADAS segmentation in HSI, so we first summarize related work on segmentation with HSI outside ADAS and deep learning-based ADAS segmentation in RGB data.

### 2.1. Deep learning-based segmentation with RGB data in ADAS

Semantic segmentation of urban scenes has been attracting the attention of researchers since the publication of widely known databases such as KITTI Stereo and Flow Benchmark [9] (200 annotated training images and 200 test images of urban environments), CamVid [10,11] (700 cityscape images), Cityscapes [12] (5000 finely-annotated and 20 000 weakly-annotated images), although there are also other databases that focus on other environments such as [13], which provides 1200 annotated images from highway scenarios. In most cases, these databases are used to compare different models according to various metrics that are usually focused on accuracy results without paying attention to either network complexity or energy consumption, neglecting their possible deployment in an embedded application.

In fact, when analyzing the architectures that report remarkable results in the Cityscapes database, it can be seen that most of them contain tens of millions of parameters and require billions of floating point operations (GFLOPS) at inference. For instance, Deeplab [14], a deep convolutional neural network which includes atrous convolution for dense feature extraction and field-of-view enlargement, contains tens of millions of parameters. SwiftNet [15], another deep convolutional encoder–decoder network which is based on an interleaved pyramid fusion model, contains, depending on the configuration, from 2.4M to 24.7M parameters and performs from 41 to 218 GFLOPs when run on an NVIDIA GTX 1080Ti GPU. STDC [16], a convolutional model which features various short-term dense concatenate modules (STDC) to fuse outputs from convolutions with different receptive fields and ends with a fully connected layer, contains, depending on the architecture, from 8.44M to 813M parameters and performs from 0.01 to 1.45 GFLOPs when run on an NVIDIA GTX 1080Ti. Finally, TinyHMSeg [17] is an encoder–decoder network which instantiates MobileNetV2 as the lightweight backbone and step-wisely reuses the intermediate feature maps, downsamples them and feeds them to medium- and low-resolution branches. Unlike those mentioned above, TinyHMSeg is the only one that aims to achieve high performance results while remaining lightweight; it contains 0.7M parameters and performs 3 GFLOPs when run on an NVIDIA GTX 1080Ti for real-time applications.

Finally, it should be noted that there is a line of research discussed in [18] that promotes precisely latency-aware, real-time video segmentation for which it seeks to leverage the temporal correlation of consecutive frames in a video to improve the next-frame prediction and reduce computation burden and latency.

### 2.2. RGB versus HSI for image segmentation

The hyperspectral reflectance data of different material surfaces outside the visible spectrum can provide very valuable information to improve intelligent vision tasks and, particularly, for image segmentation. Outside ADAS, HSI is being applied in various fields such as remote sensing for geoscience applications, food assessment and biomedical image analysis. With regards to remote sensing, in [19], to identify tree species in mixed-conifer forests, high spatial resolution airborne hyperspectral imagery and a convolutional neural network are

used. The most relevant aspect of this article is that the authors show that the HSI CNN model outperformed the RGB CNN model by 23% on average, taking into account all metrics and species. This indicates that the additional spectral information provided by HSI is essential for increasing model performance. In the field of food assessment, in [20], they show the greater potential of HSI compared with conventional RGB imaging for predicting L-value of mushrooms, a method for mushroom quality grading. Finally, concerning biomedical image analysis, especially interesting is the work presented in [21] where they compare five deep learning models, varying both the spatial (from pixel-based segmentation to image-based segmentation) and spectral (RGB, domain-specific tissue parameter images (TPI) and HSI) granularities. They show that unprocessed HSI data (not TPI) offers a huge benefit compared to RGB or TPI data for organ segmentation, although this superiority is reduced when using more complex networks such as FCNs, which take advantage not only of the richness of the spectral signature of the materials but also of the spatial properties of the images.

### 2.3. Deep learning-based segmentation with HSI data in ADAS

The main hypothesis of this line of research is that spectral information beyond the visible range can improve the accuracy and robustness of current RGB-based ADAS for object identification and tracking and, eventually, can be used for scene understanding. Leaving aside the studies carried out with thermal cameras (far infrared), there is not much previous work published in the field [22]. Nevertheless, there are some researchers that have started to explore the applicability of HSI cameras to the field of ADS. In [23], a hyperspectral image dataset for road segmentation in urban and rural scenes is introduced and in [7] the authors explore the use of FCNs to segment HSI applied to ADAS. The underlying idea is that the incorporation of richer spectral information can provide a distinct spectral fingerprint for each entity in the image, helping to achieve more precise and robust detection systems.

In the study conducted by [2], the authors characterize the reflection of a wide variety of road objects, such as a work zone signaling cone, a road reflector, an asphalt pavement, a dark blue fleece sweater, a tyre and an exhaust pipe in the 350 to 2450 nm range using a spectroradiometer. They show that there is a great difference among the spectral signatures of all the measured materials and remark that the zone of greatest reflectance is between 750 and 1550 nm (NIR-SWIR bands). In the same way, in [5] they evaluate the relevance of four different spectral bands: RGB, NIR, SWIR and LWIR when detecting and recognizing pedestrians, vehicles, traffic signs and lanes both in outdoor natural conditions and in an artificial fog tunnel. They conclude that, overall, the NIR band is the most useful spectral band followed by SWIR, RGB and LWIR. Another work which follows the same line is [3], where the authors generate realistic spectral radiance data of synthetic scenes in the VIS and NIR spectral range and assess the benefits of VIS-NIR imaging in their synthetic images using exemplary night time and daylight traffic scenes. Their results for daylight scenes confirm the benefits of using the additional information provided by the NIR spectral band for metamerism resolution, although they acknowledge that the nighttime scene results are not conclusive and more models of artificial light sources should be added to the database. In [24], based on the uniqueness of the spectral signature of different materials, the authors propose a pixel-wise labeling process aided by hyperspectral imaging which can, to a great extent, reduce the manual labeling process. Another paper that deals with the metamerism phenomenon is [4], where they introduce their hyperspectral dataset for semantic segmentation of cityscapes and remark that metamerism is particularly challenging in these scenes because they contain too many classes, complex lighting and spatial structures. They perform tSNE visualization of RGB and HSI data, and show that HSI images have a stronger separability. One of the earliest studies in this regard, although



(a) Ground truth image.

(b) False-color visible image.

**Fig. 1.** Ground truth image (a) and its corresponding false-color visible image (b) of an urban scenario as an example of the weak labeling methodology followed to develop the HSI-Drive dataset. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

not specifically aimed at ADAS, draws interesting conclusions regarding the ability of images taken with hyperspectral cameras to detect people in urban environments [25]. Quite interestingly, one of the outcomes of the work is the finding that the use of the VNIR spectrum (up to 1000 nm) offers similar results to using the full spectrum of the used sensor (up to 2500 nm). This is very relevant to the present proposal, since current CMOS sensor-based low-cost HSI snapshot cameras do not offer spectrometric information beyond 1000 nm.

Another initial hypothesis of this research consists of the assumption that deep-learning algorithms developed for image segmentation in the field of ADAS will take advantage of hyperspectral data, either to get lighter network models for the same task (since the input information is richer, it is expected that the more informative HSI images will require a less complex processing pipeline and so, they will be more suitable for embedded applications) or by improving the performance obtained with RGB images. [26] is a paper that focuses exactly on analyzing these two aspects for a video-surveillance system. They show that while using a pixel-wise classification neural network, the error rate decreases by 30% when using RGB-NIR data together compared to using RGB data alone. In addition, they take into account the hard power constraints of embedded platforms and train the RGB-NIR data on two convolutional networks with smaller computational effort, getting the same error rate as the RGB-only pixel-wise neural network, while decreasing the computational load by 300%. A pioneering research group in the investigation of the use of HSI in ADS is the Active Vision Group (AGAS) of the University of Koblenz-Landau. They have published various papers on this topic, reporting interesting results on image segmentation and terrain classification applied to images combining VIS and NIR spectrum information from HSI snapshot cameras. In [27] the authors satisfactorily combine a 16-band image in the visible spectrum with a 25-band image in the NIR spectrum obtaining promising results and improve the performance of a method which only relies on a RGB image and a 1-band NIR image. In [28], the work is focused on extracting robust spatial and spectral features for snapshot hyperspectral data which are invariant to illumination changes and then train a per-pixel classification network. In [29], the authors propose a snapshot hyperspectral scene analysis pipeline which combines a per-pixel classification with context-aware, fully connected, conditional random fields leading to a more accurate pixel-level classification performance. Of special interest is [6], where the authors use a deep autoencoder with additional regularization terms to learn the latent space of the input hyperspectral snapshot data and obtain a reduced feature set which, when trained with deep learning methods, outperforms the RGB data networks.

In summary, prior work on semantic segmentation with real HSI data from snapshot cameras in ADAS is scarce and, in general, little attention has been paid to the realistic deployment of these system and to their characterization in terms of processing performance. This includes analyzing not only the segmentation process itself, but also the data acquisition and the raw image preprocessing stages. There

are some works that use synthetic databases that are not capable of accurately reproducing the complex illumination of real traffic scenarios [3]. Many papers are aimed at detecting or identifying road objects but not at performing semantic segmentation of images for scene understanding [2,5]. The few works that combine real data with image segmentation tasks do not address the problem of analyzing how these systems can be realistically deployed on a resource and power constrained embedded computing platform while keeping processing latency values in the range of what is acceptable for ADAS [4,6,27–29].

### 3. Experimental setup

#### 3.1. The dataset

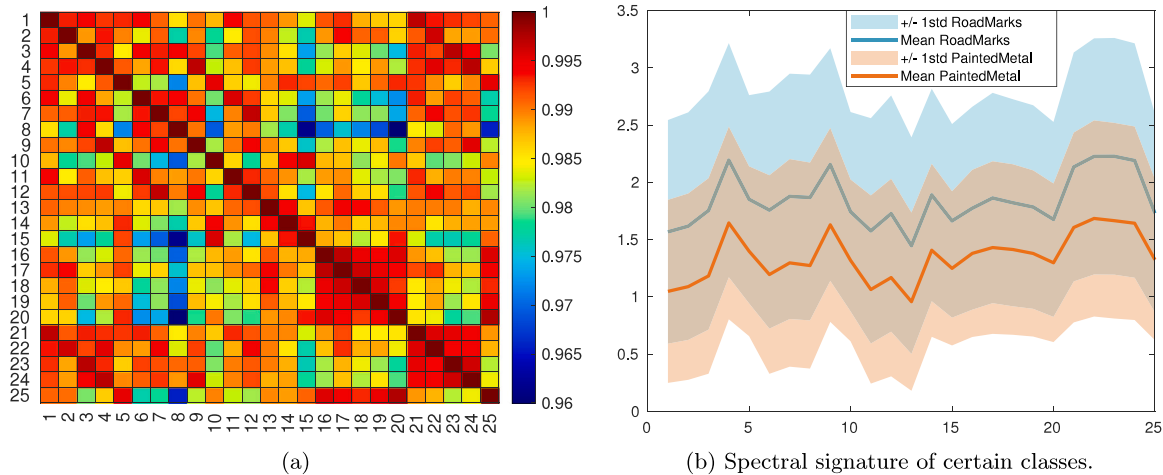
As is reported in [8], there are very few HSI datasets aimed at developing ADAS and ADS applications. In [8], the authors present a dataset named HSI-Drive, which was specifically conceived to provide the research community with data obtained at real driving scenarios with a small size, portable, snapshot hyperspectral camera capable of recording 25-band HSI at video rates. The current version of the dataset, i.e. HSI Drive v1.1, contains 276 manually labeled images recorded while driving a car in urban, road and highway scenarios under diverse weather (sunny, cloudy, rainy and foggy) and lighting (dawn, midday, sunset) conditions. The recordings were performed during Spring (121 images) and Summer (155 images). The recording setup included a single Photonfocus camera featuring an Imec 25-band VIS-NIR (535 nm–975 nm) mosaic spectral filter over a CMOSIS CMV200 image wafer sensor. The resolution of the raw images was  $1088 \times 2048$  pixels with  $5 \mu\text{m} \times 5 \mu\text{m}$  size. However, as the spectral bands were extracted from a mosaic formed by  $5 \times 5$  pixel window Fabri-Perot filters, the output resolution of the HSI cubes was actually reduced to  $216 \times 409 \times 25$  [30]. Indeed, the use of this technology implies performing an image preprocessing stage prior to the processing pipeline to get the spectral cubes from the recorded radiance raw images, an issue that is thoroughly addressed in Section 5.2 of this paper.

The original labeling of the scene images in the dataset comprises 10 different classes related mostly to the expected differences in the surface reflectance signatures of various materials: metal, vegetation, concrete, tarmac etc. Available ground truth masks however are not fully dense, since the labeling method was aimed at favoring the accuracy in the spectral information provided for each class by reducing spectral mixing. This means, for instance, that contour pixels of many items and the surfaces on the background that cannot be precisely distinguished have not been labeled (Fig. 1). While this method of labeling is not an issue for non-convolutional ML models, it may have some adverse effects on the performance of CNN-based segmentation models. In fact, the training of convolutional neural networks with “weak supervision” on sparsely labeled datasets is a line of research itself (see for instance [31] and references therein).

**Table 1**

Jeffreys-Matsushita (JM) interclass distances which range from 0 (null separability) to 2.0 (perfect separability).

	Road	Road M.	Veg.	P. Met.	Sky	Conc.	Ped.	Unp. Met.	Glass
Road		1.92	1.83	1.65	1.98	1.44	1.84	1.42	1.49
Road marks	1.92		1.79	1.63	1.87	1.68	1.92	1.92	1.93
Vegetation	1.83	1.79		1.44	1.96	1.64	1.81	1.73	1.81
Painted metal	1.65	1.63	1.44		1.91	1.48	1.70	1.35	1.48
Sky	1.98	1.87	1.96	1.91		1.97	1.98	1.98	1.80
Concrete	1.44	1.68	1.64	1.48	1.97		1.74	1.62	1.68
Pedestrian	1.84	1.92	1.81	1.70	1.98	1.74		1.79	1.66
Unpainted metal	1.42	1.92	1.73	1.35	1.98	1.62	1.79		1.38
Glass	1.49	1.93	1.81	1.48	1.80	1.68	1.66	1.38	
<b>Mean</b>	<b>1.73</b>	<b>1.80</b>	<b>1.78</b>	<b>1.63</b>	<b>1.96</b>	<b>1.69</b>	<b>1.82</b>	<b>1.68</b>	<b>1.71</b>



**Fig. 2.** Pearson Correlation Coefficients among the 25 spectral bands (a) and spectral signature of some of the representative classes (b) of the HSI-Drive 1.1 dataset.

Regarding the spectral characteristics of the data, in [7] the authors provide spectral separability figures of all 10 classes calculated on all available data in the HSI-Drive v1.1 dataset. In particular, they use the Jeffreys-Matsushita (JM) index, which ranges from 0 (null separability) to 2.0 (perfect separability) as the preferred criterion, which is relevant in this work since this index provides an estimation of the probability of a correct prediction to happen [32]. As can be observed in Table 1, which we reproduce here for clarity, the separability of some of the class pairs such as Road/Road Marks (1.92), Road/Sky (1.98) or Road Marks/Unpainted Metal (1.92) is promisingly high, while class-pairs such as Road/Concrete (1.44), Road Marks/Painted Metal (1.63) and Painted Metal/Unpainted Metal (1.35) show low separability indexes. In general, JM index values below 1.9 imply that designing classification algorithms based on pure spectral reflectance data may be challenging or infeasible, and that more meaningful and relevant features should be extracted from the raw data to achieve good performance in an image segmentation task.

Beyond possible similarities of the “real” physical spectral reflectance footprints of different material surfaces in the spectral range of the sensor, low separability indexes between categories in the dataset can be a consequence of three main factors: redundant information in the spectral dimension, high variability within the same category due to changing lighting and different camera set-ups, and the inherent technological limitations of mosaic-filters to accurately separate the spectral radiance that enters the sensor. Examples of this are the presence of crosstalks, the variability produced by different light beam angles of incidence and the spectral mixing produced in the demosaicing process.

The aforementioned limitations associated with snapshot technology lead to a correlation matrix (Fig. 2(a)) which shows a strong Pearson Correlation Coefficient among the bands.

In order to determine the existing intraclass variability, we have used our own spectral library. This library consists of 10 images (one for each class) which have been taken on the same day, at the same

time and with the same camera configuration. The objective is to ascertain what the properties of the materials are in the most controlled conditions as possible (while still acquiring them outdoors with the snapshot hyperspectral camera). Fig. 3 shows the maximum intrinsic intraclass Pearson Correlation Coefficients of two representative classes of the dataset. It can be seen that although Painted Metal is quite homogeneous, the dispersion of Road Marks is heavily affected by the nature of the materials.

Finally, Fig. 2(b) displays the area covered by the mean  $\pm 1$ std of the spectral signatures of the same two classes, using more than 250,000 pixels for each class. The Figure confirms that interclass variance is high and results in a wide overlapping among spectral signatures which, on the other hand, is something common in other HSI datasets (see [33] for instance).

Despite the issues described above and the challenges posed by the use of data acquired in real outdoor driving scenarios, as mentioned in Section 2, there are numerous works that confirm the benefits of using HSI in combination (or not) with RGB over using only RGB images for image segmentation. Since the camera used to develop the dataset does not cover the entire visible spectral range, it is not possible to generate equivalent RGB images to perform a proper comparison of classification performance to standard RGB. It is possible however to select the three most informative bands out of the original 25 in the dataset and generate pseudoRGB images to perform meaningful comparative analyses. In particular, we used the orthogonal space projection method proposed in [34] to identify the most informative spectral bands from every image in the dataset and then selected the most frequent (mode) values, which happened to be the triplet (8, 21, 24) in the mosaic filter. Applying a t-Distributed Stochastic Neighbor Embedding (t-SNE) algorithm [35], it is possible to map the data on a 2D to graphically visualize the separability of the classes. Fig. 4(a) shows the result of applying the t-SNE to the pseudoRGB images, while in Fig. 4(b), the results when using all 25 bands are shown. It can be seen that when




Fig. 3. Intraclass Pearson Correlation Coefficients for the 25 bands of two representative classes of the dataset: Road Marks (a) and Painted Metal (b).

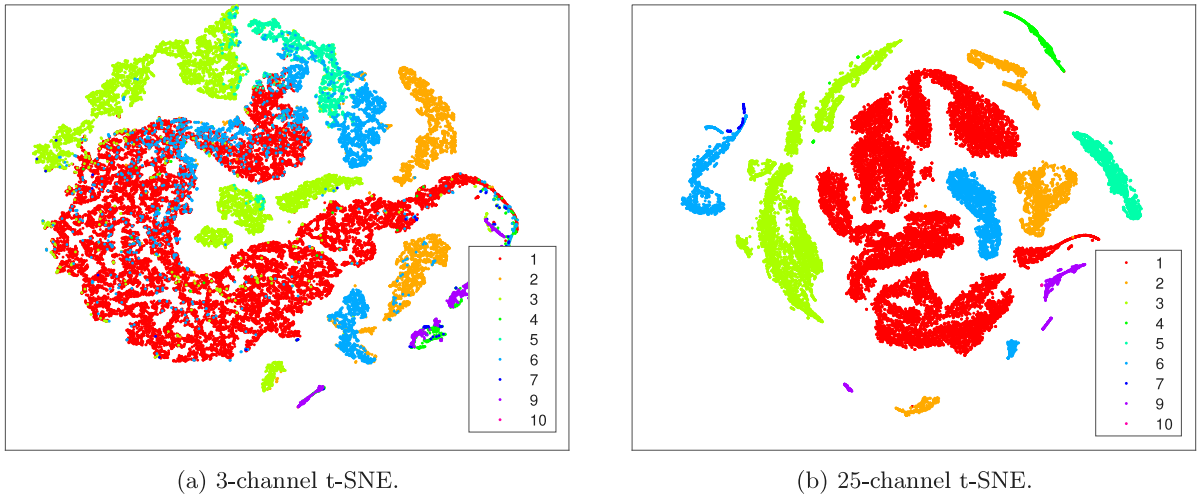


Fig. 4. Comparison of t-SNE output when using either the 3 most representative channels (a) or the full 25 channels (b).

using all bands, the data representation is not only free of overlapping among classes, but is also more compact, as clusters of the same class are closer together. The contribution of the hyperspectral information compared to pseudoRGB data is further analyzed in Section 3.3 by evaluating the performance of baseline spectral classifiers.

### 3.2. Data partitioning and application examples

In order to investigate the potential performance of segmentation systems based on the training of ML models with this dataset, throughout this work we use two application examples that can be employed in the development of ADAS/ADS. The first example focuses on segmenting images into three classes: road (tarmac), road marks and non-drivable areas, the latter including the remaining classes in the labeled masks. This low-complexity segmentation example would be aimed at developing a system for the discrimination of drivable and non-drivable zones, together with a lane-keeping aid for ADAS or even a trajectory planning system for ADS. In the second example, two additional classes have been added to the model training (vegetation and sky), which show quite satisfactory spectral separability indexes (see Table 1) and add relevant information for a scene understanding application. The exploration of more complex segmentation models including all 10 classes in the dataset has also been investigated and

some segmented images are shown in Section 4. However, the results obtained are irregular and inconclusive, and will require further research.

In order to perform an ML training on this dataset, the 276 available images have been divided into training, validation and testing subsets. The selection of images has not been aleatory but conducted following a data-diversity criterion. The generated subsets contain 162 images for training, 57 for validation and 57 for testing, preserving class proportionality in all three subsets. As for the metrics used in the evaluation of the performance tests, these have been Recall (R), Precision (P), and Intersection over union (IoU), which is an equivalent for the IoU for object detection. As Eqs. (1)–(3) show, recall accounts for the false negatives (FN) and true positives (TP), precision takes into account the false positives (FP) and the true positives too and IoU combines both aspects:

$$R_i = \frac{TP_i}{TP_i + FN_i} \quad WR = \sum_{i=1}^n w_i R_i \quad (1)$$

$$P_i = \frac{TP_i}{TP_i + FP_i} \quad WP = \sum_{i=1}^n w_i P_i \quad (2)$$

$$IoU_i = \frac{TP_i}{TP_i + FN_i + FP_i} \quad WIoU = \sum_{i=1}^n w_i IoU_i \quad (3)$$

**Table 2**  
Frequency of each of the classes in the train/val and test subsets for the 3-class and 5-class experiments.

		Total	3-class experiment			5-class experiment				
			Road	Marks	No Dri.	Road	Marks	Veg.	Sky	Other
Train	Num. pixels	13 343 314	7 947 254	450 703	4 945 627	7 947 254	450 703	3 127 191	568 411	1 250 025
	%	100	59.56	3.38	37.06	59.56	3.38	23.44	4.26	9.37
		Total	Road	Marks	No Dri.	Road	Marks	Veg.	Sky	Other
Test	Num. pixels	3 514 081	2 067 379	99 426	1 347 276	2 067 379	99 426	820 804	163 127	363 345
	%	100	58.83	2.83	38.34	58.83	2.83	23.36	4.64	10.34

where  $i$  is the class index such that, for example,  $FN_i$  accounts for the pixels that have been predicted as not belonging to class  $i$ , but are actually part of class  $i$ . As a consequence of the dataset being heavily imbalanced (Table 2), a correct interpretation of the indexes requires not only a representation of the global metrics, but also the mean values and, more specifically, the weighted scores of performance indexes. Accordingly, weighting factors related to the inverse of the frequency of the classes in the dataset have been added to the final calculation of the indexes.

### 3.3. Reference baseline spectral classifiers

In a first approach to the problem, baseline neural classifiers were trained relying only on spectral data. The results obtained, in addition to exposing the foreseeable limitations of segmentation based on pure, non-processed spectral features, will be useful for assessing fairly the performance results of the convolutional models presented in the following sections.

We trained medium-depth, fully connected ANNs with three hidden layers to perform image segmentation on the per-pixel classification basis. First, spectral reflectance values were normalized through division by the sum of the reflectance at all 25 wavelengths to alleviate the adverse effects of differing illumination at the image level. The models featured Hyperbolic Tangent activation functions at every layer and included an input z-score normalization layer and a Softmax output layer. Due to large dataset size and memory limits in the NVIDIA GFORCE RTX-3090 used for training (24 GB), the optimization was performed by an Adam (Adaptive Moment estimation) algorithm with  $2^{20}$  mini-batch size, data shuffling at every epoch and a limit of 300 epochs. Selected network parameter settings corresponded to the best cross-entropy loss value in the validation set. Due to severe class imbalance in the training dataset, various techniques such as under-sampling to the smallest class dataset size, data augmentation by SMOTE algorithm, weighting of the loss function and changing to a Focal Loss function were tried out. None of these techniques produced entirely satisfactory results, we thus applied a combination of techniques by generating some synthetic data to reduce class imbalance. Then, a weighted, focal loss function was set, based on data frequency. Nevertheless, the weights were finally manually adjusted until satisfactory results were achieved.

After various attempts, ic-25-100-100-c ANN models, where ic stands for the number of input channels and c for the number of classes to be categorized, showed good performance, while adding more nodes/parameters did not produce any accuracy improvements. Table 3 sums up the results obtained for the two example applications under study. While performance figures are generally acceptable, when per image indexes are analyzed, there is a lack of robustness in the predictions. In particular, the precision of minority classes tends to be low, with a high rate of false positives in some images, which is reflected in the lower weighted scores. Figs. 10 and 11 show produced segmentation for three example images in the testing set, one for each type of driving scenario: urban, road and highway. As can be seen, produced segmented images are overall correct, producing interpretable scenarios, but the lack of precision, which produces incorrectly classified pixel infiltration, mostly from the majority class into minority class segments, is noticeable.




Fig. 5. Overall, average and weighted IoU (%) as a function of the number of spectral channels.

The same experiment was performed using pseudoRGB images generated with the three most VNIR informative bands for each application example. After various training setting adjustments, obtained best results are summed up in Table 3. As can be observed, the classifiers are unable to generate meaningful results, thus confirming, as expected, the impossibility of designing an image segmentation system based only on spectral reflectance data with just three channels, and demonstrating the superiority of HSI in this task. Additionally, the ANN model has been iteratively trained for an increasing number of spectral band data in the input vectors, ranging from just 1 band to all 25 bands. Fig. 5 shows obtained performance scores on the testing set for the first segmentation example (3 classes). It can be noted that, despite the strong spectral cross-correlation of bands, the accuracy of the classifier improves as more spectral information is provided to the input which, in principle, discourages performing spectral channel selection to reduce complexity.

### 4. FCNs for HSI image segmentation

It is well known that incorporating spatial information to the spectral data in HSI can be an effective means to overcome the limits of the spectral imaging technology for image segmentation. The use of convolutional filters with tunable parameters in CNNs to extract spatial information has shown excellent performance in many applications. In particular, FCNs are neural computing architectures specifically aimed at image segmentation. However, when convolutional models are designed and trained with disregard for the computational constraints imposed by the target processing devices and the requirements of the application (power, latency), the results obtained may prove to be impractical. In this section we describe the development of an FCN that enhances the accuracy and robustness of the spectral classifiers while keeping model architecture in the range of tens of thousands of parameters for efficient embedded processing.

**Table 3**  
Evaluation of the ANN on the 3-classes (above-top) and 5-classes (above) test datasets in terms of recall, precision and IoU (Intersection over Union).

	25 channels			pseudoRGB		
	Recall	Precision	IoU	Recall	Precision	IoU
Road	89.39	92.91	83.67	81.70	78.37	66.66
Road marks	48.88	29.66	22.64	04.47	19.93	03.79
Non-drivable	90.34	89.42	81.62	67.94	68.48	51.75
Overall	<b>88.61</b>	<b>89.78</b>	<b>81.16</b>	<b>74.24</b>	<b>72.92</b>	<b>59.17</b>
Mean	<b>76.21</b>	<b>70.66</b>	<b>62.64</b>	<b>51.37</b>	<b>55.59</b>	<b>40.73</b>
Weighted	<b>53.35</b>	<b>36.30</b>	<b>29.13</b>	<b>11.95</b>	<b>25.63</b>	<b>09.64</b>
Road	85.71	94.18	81.39	59.91	87.22	55.07
Road marks	54.34	26.85	21.91	38.80	08.32	07.36
Vegetation	95.05	90.29	86.24	75.95	77.45	62.20
Sky	96.05	82.76	80.36	69.21	30.27	26.68
Other	56.72	54.52	38.50	24.85	19.97	12.45
Overall	<b>84.49</b>	<b>86.74</b>	<b>76.35</b>	<b>59.87</b>	<b>73.11</b>	<b>49.66</b>
Mean	95.85	69.719	68.84	53.75	44.65	32.75
Weighted	<b>70.16</b>	<b>52.45</b>	<b>46.55</b>	<b>48.64</b>	<b>22.32</b>	<b>18.12</b>




Fig. 6. Architecture of the modified U-Net.

#### 4.1. Model development

U-Net is an FCN-based architecture intended for accurate image segmentation. This model is of the type encoder–decoder, which means that after an encoding stage, there is a decoding or “deconvolution” up-sampling stage that re-projects the extracted characteristics, and recovers the original image resolution to directly produce segmentation maps. Intermediate skip connections enable the fusion of low- and high-level features. U-Net was originally intended for biological image segmentation [36] but has been successfully used in other segmentation tasks, such as precision agriculture [31], food quality assessment [37] and aerial city recognition [38]. The idea of using an FCN to process HSI is to combine the intrinsic spectral characteristics of the different classes with the spatial relationships extracted by the convolution operations.

We have tailored the original architecture of U-Net to the unique characteristics of the dataset (see Fig. 6, made with PlotNeuralNet software [39]) by including batch-normalization layers after convolutional filters, adding a dropout layer before the decoder, adapting the padding of the convolutional filters and modifying the network depth and the number of filters in the initial convolutional block. Besides this, we

have also customized some training-related parameters which will be discussed when the training procedure is described.

In addition to what has been explained in the previous paragraph, to get the best trade-off between segmentation performance and computational complexity, a hyperparameter optimization process has been carried out by programming a grid search of the optimum combination of a set of model hyperparameters. The optimization performance has been evaluated by calculating the segmentation accuracy on a subset of 45 images selected from all possible environments and weather conditions. After an initial evaluation of the model’s performance, the selected set of analyzed hyperparameters included: the size of the input image patches [64 × 64, 128 × 128], the encoder depth [2, 3, 4] and the number of filters in the first convolutional block [8, 16, 32]. Additionally, the optimization procedure also included the influence of two parameters related to the training process: the initial learning rate of the optimization algorithm (Adam) and the stride of the patch creation process. Regarding the initial learning rate, we opted for varying it in decreasing powers of 10 from 10<sup>-1</sup> to 10<sup>-4</sup>. As for the stride, we have opted for training the FCN with patches and not with the full resolution image so as to, on the one hand, have a lighter neural network and, on the other hand, benefit from the fact that as






Fig. 7. Example of the stride selection effect in the overlapping matrix index using 18 (3 × 6) patches.

some of the patches overlap in certain pixels, there is more information to correctly predict the class to which those pixels belong. How the patches overlap depends both on the starting position of the initial patch (upper left side of the image in our case) and on the value given to the horizontal and vertical strides. We have decided that the areas where we want the segmentation to be as accurate as possible (and consequently locate the most overlapping patches) should be the central area of the image (a centrosymmetric overlapping scheme as Fig. 7 shows) and accordingly set the vertical and horizontal stride values.

The overlapping matrix is defined in Eq. (4), which is simply the sum of the patch masks (binary matrices of the same size as the input image containing ones where the patch is defined and zeros everywhere else).

$$OI = \sum_{i=1}^N pm_i \in \mathbb{R}^3 \quad (4)$$

where  $N$  is the number of patches and  $pm_i$  is the patch mask associated to patch number  $i$ . The results from the hyperparameter tuning process are gathered in Tables 4 and 5. It can be seen that using a larger 128 patch-size tends to produce a lower validation error than when using the smaller 64 patch-size. With regard to the number of initial filters and the encoder depth there is not a clear pattern and the results depend on the particular combination of different parameter values.

In order to establish a numerical criterion for the selection of the best model, we have evaluated two metrics that give an insight into the memory footprint and the computing complexity of the model architecture: the total number of parameters (NP) and the number of floating point operations (FLOPS). Both metrics depend on some other variables that were frozen as we did not take them into account in the hyperparameter tuning study. More explicitly, the number of parameters, as shown in Eq. (5), is a function of the encoder depth ( $ed$ ), the initial number of filters ( $if$ ), the convolution/up-convolution kernel sizes ( $cks$ ,  $ucls$ ), the number of classes to be predicted ( $c$ ) and the number of input image channels ( $ic$ ). The number of FLOPS (Eq. (6)) depends on the preceding parameters as well as on the input image spatial dimensions ( $id$ ) and other architectural choices such as the strides of the convolution/deconvolution and pooling layers.

$$NP = f(ed, if, cks, ucls, c, ic) \Big|_{cks=3, ucls=2, c=3/5, ic=25} \quad (5)$$

$$FLOPS = f(ed, if, cks, ucls, c, ic, id) \Big|_{cks=3, ucls=2, c=3/5, ic=25, id=128} \quad (6)$$

Fig. 8 shows the evaluation of the NP and FLOPS in the range of the hyperparameters under analysis. It can be seen, on the one hand, that




Fig. 8. Variation in the number of FLOPS and parameters as a function of ( $ed$ ,  $if$ ).




Fig. 9. Best validation accuracy in terms of encoder depth and initial number of filters.

the NP is almost quadrupled as we increase either the encoder depth or the initial number of filters. On the other hand, the number of FLOPS significantly increases with the initial number of filters, and not quite with the growth of the encoder depth.

Fig. 9 depicts the highest achieved validation accuracy for each pair ( $ed$ ,  $if$ ) after training. Increasing  $ed$  slightly improves the validation accuracy regardless of the initial number of filters.

All in all, combining the information extracted from both graphics, we selected the combination ( $id$ ,  $ed$ ,  $if$ ) = (128, 2, 8) as the architecture with the best trade-off between model complexity and classification performance. Not only does it perform very similarly to the other architectures in terms of segmentation, but choosing a higher number of initial filters would considerably increase the memory footprint and the number of FLOPS (without improving precision/recall), and increasing the model depth would significantly increase the ratio of memory accesses with respect to the number of FLOPS (overshadowing the minor increase in precision/recall).

We can now compare the U-Net and the baseline ANN in terms of model size and computational complexity (MACS, Multiply and Accumulate operations or FLOPS) as it is reflected on Table 6: while the ANN has only 13,653 parameters and performs 1,206,000,000 FLOPS (27,093 FLOPS per pixel) during inference, U-Net model has 31,707 parameters (320 of them are non-trainable) and needs 4,678,200,00 FLOPS (259,900,000 FLOPS per patch) to produce an output. The difference between the FLOPS ratio (3.879x) and the parameters ratio (2.3x) affects the time needed to make the forward pass, which will be assessed in the next section.

#### 4.2. Segmentation results

Having first selected the best model architecture, we have trained U-Net on the complete training set. The training has been performed on a NVIDIA GeForce RTX 3090-24 GB GPU with an Adam optimizer (0.9 gradient decay factor and 0.999 squared gradient decay factor), a

**Table 4**

Hyperparameter analysis for the following configurations: patch size, 64; initial number of filters, 8, 16, 32; encoder depth, 2, 3, 4 and initial learning rate, 0.1, 0.01, 0.001 and 0.0001.

Initial num. of filt.		8			16			32		
Enc. depth	Img. size	2	3	4	2	3	4	2	3	4
		<b>64</b>	<b>0,100</b>							
<i>maxValidationAccuracy</i>		90.55	90.48	92.74	91.11	90.97	93.29	92.56	94.49	94.52
<i>minValidationLoss</i>		0.28	0.35	0.23	0.24	0.28	0.20	0.25	0.17	0.20
<i>maxTrainingAccuracy</i>		94.34	94.66	96.08	95.58	94.42	96.52	94.97	96.17	97.35
<i>minTrainingLoss</i>		0.21	0.17	0.13	0.16	0.14	0.11	0.15	0.12	0.07
<b>0,01</b>										
<i>maxValidationAccuracy</i>		<b>94.65</b>	<b>96.67</b>	<b>96.48</b>	94.79	96.41	<b>96.34</b>	94.94	96.32	94.80
<i>minValidationLoss</i>		<b>0.16</b>	<b>0.13</b>	<b>0.16</b>	0.18	0.14	<b>0.12</b>	0.17	0.14	0.20
<i>maxTrainingAccuracy</i>		<b>97.50</b>	<b>98.64</b>	<b>98.66</b>	96.70	99.16	<b>98.88</b>	97.67	98.95	97.97
<i>minTrainingLoss</i>		<b>0.08</b>	<b>0.04</b>	<b>0.04</b>	0.10	0.03	<b>0.04</b>	0.06	0.03	0.06
<b>0,001</b>										
<i>maxValidationAccuracy</i>		93.69	95.09	95.40	<b>95.69</b>	<b>97.12</b>	96.02	<b>96.31</b>	<b>96.99</b>	<b>95.89</b>
<i>minValidationLoss</i>		0.21	0.17	0.20	<b>0.19</b>	<b>0.13</b>	0.16	<b>0.16</b>	<b>0.14</b>	<b>0.17</b>
<i>maxTrainingAccuracy</i>		95.10	98.01	98.85	<b>98.22</b>	<b>99.08</b>	99.40	<b>98.23</b>	<b>98.73</b>	<b>99.23</b>
<i>minTrainingLoss</i>		0.15	0.06	0.06	<b>0.06</b>	<b>0.03</b>	0.02	<b>0.05</b>	<b>0.04</b>	<b>0.03</b>
<b>0,0001</b>										
<i>maxValidationAccuracy</i>		86.19	79.98	90.27	86.88	92.40	92.04	93.69	95.29	94.85
<i>minValidationLoss</i>		0.43	0.67	0.37	0.37	0.27	0.27	0.26	0.19	0.21
<i>maxTrainingAccuracy</i>		88.37	84.54	93.85	89.00	94.93	97.20	94.51	97.41	98.96
<i>minTrainingLoss</i>		0.39	0.48	0.25	0.28	0.17	0.11	0.14	0.08	0.04

**Table 5**

Hyperparameter analysis for the following configurations: patch size, 128; initial number of filters, 8, 16, 32; encoder depth, 2,3,4 and initial learning rate, 0.1, 0.01, 0.001 and 0.0001.

Initial num. of filt.		8			16			32		
Enc. depth	Img. size	2	3	4	2	3	4	2	3	4
		<b>128</b>	<b>0,1</b>							
<i>maxValidationAccuracy</i>		94.31	91.65	96.41	94.10	96.37	93.60	95.48	96.23	96.32
<i>minValidationLoss</i>		0.17	0.22	0.15	0.18	0.12	0.20	0.15	0.14	0.14
<i>maxTrainingAccuracy</i>		96.91	95.11	98.56	96.91	98.89	95.98	97.74	98.46	98.33
<i>minTrainingLoss</i>		0.09	0.14	0.04	0.08	0.03	0.12	0.07	0.04	0.05
<b>0,01</b>										
<i>maxValidationAccuracy</i>		<b>95.75</b>	<b>97.25</b>	<b>97.78</b>	96.05	<b>96.64</b>	<b>97.43</b>	95.06	97.23	97.26
<i>minValidationLoss</i>		<b>0.17</b>	<b>0.13</b>	<b>0.11</b>	0.16	<b>0.16</b>	<b>0.11</b>	0.19	0.13	0.12
<i>maxTrainingAccuracy</i>		<b>98.52</b>	<b>99.49</b>	<b>99.65</b>	98.90	<b>99.49</b>	<b>99.62</b>	98.65	99.67	99.68
<i>minTrainingLoss</i>		<b>0.04</b>	<b>0.01</b>	<b>0.01</b>	0.03	<b>0.01</b>	<b>0.01</b>	0.04	0.01	0.01
<b>0,001</b>										
<i>maxValidationAccuracy</i>		95.56	97.01	96.97	<b>96.38</b>	96.57	97.02	<b>96.73</b>	<b>97.24</b>	<b>97.31</b>
<i>minValidationLoss</i>		0.17	0.16	0.17	<b>0.16</b>	0.17	0.13	<b>0.16</b>	<b>0.15</b>	<b>0.11</b>
<i>maxTrainingAccuracy</i>		97.02	99.34	99.72	<b>98.63</b>	99.62	99.77	<b>99.26</b>	<b>99.83</b>	<b>99.85</b>
<i>minTrainingLoss</i>		0.08	0.03	0.02	<b>0.04</b>	0.01	0.01	<b>0.02</b>	<b>0.01</b>	<b>0.01</b>
<b>0,0001</b>										
<i>maxValidationAccuracy</i>		73.08	91.65	93.33	93.41	94.96	94.23	93.88	95.56	96.19
<i>minValidationLoss</i>		0.64	0.31	0.25	0.24	0.24	0.21	0.29	0.20	0.19
<i>maxTrainingAccuracy</i>		81.81	94.12	97.20	95.28	97.95	99.07	97.19	99.10	99.77
<i>minTrainingLoss</i>		0.50	0.31	0.10	0.16	0.08	0.03	0.09	0.02	0.01

**Table 6**

Comparison of U-Net and ANN with regard to model size and computational complexity.

	U-Net	ANN
Num. of params.	31.707 K	13.653 K
FLOPS	259.9 M	27.093 K
FLOPS per image inference	4.678 G	1.206 G

mini batch size of 128, an initial learning rate of 0.005 with drop period and weighted cross-entropy loss function, among other parameters, for, at most, 60 epochs. Table 7 collects the classification performance on the patches in the testing set. In addition, as previously stated, the images are reconstructed to the original resolution from the 18 overlapping patches into which each of them has been divided, so Table 7 also displays the metrics for the complete reconstructed images. The comparison shows how the use of overlapping patches improves the segmentation performance, especially the precision, compared to the non overlapping reconstruction. This is because the FCNs tend to

fail to correctly predict the classes on the patch contours since they lack part of the contextual information for the contour pixels.


Analyzing the numerical results of the reconstructed images, it can be seen that all the classes show a good IoU value with the exception of the Road Marks class, which suffers from a low precision value. In the first example, in particular, for every 100 TPs of this class there are 37 FPs. However, as Road Marks is the minority class, this inaccuracy does not severally affect the overall metrics. This can be visually verified when depicting the segmentation outputs for the testing images. Figs. 10 and 11 show the segmentation maps of three example images chosen from the three different driving scenarios in the dataset: an urban environment, an interurban road and a highway. The proposed FCN quite satisfactorily segments a typical driving scene for the 3-class experiment (Fig. 10(k)) while it misinterprets some pixels in more challenging images such as those where there are objects casting their shadows on the road (Fig. 10(j)) or where there are overlapping objects, as on the left side in (Fig. 10(l)).

Similarly, in the five-class example (Figs. 11(j) to 11(l)), similar issues can be observed, but there is much more information about the

**Table 7**

Performance of the modified U-Net (patches and rebuilt images from overlapping patches) on the 3-classes (up) and 5-classes (down) test datasets in terms of recall, accuracy and IoU (Intersection over Union).

	Patches (128 × 128 × 25)			Rebuilt images from overlapping patches		
	Recall	Precision	IoU	Recall	Precision	IoU
Road	97.90	95.66	93.74	98.54	94.56	93.25
Road marks	90.25	73.11	67.75	87.89	77.22	<b>69.80</b>
Non-drivable	91.07	97.16	88.71	91.20	98.57	<b>90.01</b>
Overall	95.37	95.55	91.31	95.42	95.44	91.50
Mean	93.07	88.64	83.40	92.54	90.12	84.35
Weighted	90.60	75.71	70.27	88.54	79.43	<b>72.60</b>
Road	92.61	99.05	91.36	93.28	99.00	<b>92.41</b>
Road marks	80.93	75.39	64.02	78.32	79.11	64.90
Vegetation	94.98	94.63	90.12	95.74	95.80	<b>91.88</b>
Sky	97.86	93.09	91.23	97.49	93.39	91.20
Other	84.97	62.71	56.45	84.83	64.59	<b>57.90</b>
Overall	91.79	93.29	86.47	92.75	93.80	<b>87.66</b>
Mean	90.18	84.98	78.64	89.93	86.38	<b>79.66</b>
Weighted	88.64	82.14	75.27	87.40	84.15	75.93



**Fig. 10.** Comparison among the visible (first row), ground truth (second row), ANN segmentation (third row), U-Net segmentation (fourth row) images of three different scenarios: urban (first column), road (second column) and highway (third column) for the 3-class experiment.




Fig. 11. Comparison among the visible (first row), ground truth (second row), ANN segmentation (third row), U-Net segmentation (fourth row) images of three different scenarios: urban (first column), road (second column) and highway (third column) for the 5-class experiment.

scenes. For instance, it can be observed that the system is now able to identify the presence of some objects in the non-drivable sections of the images such as traffic signals, pedestrians and guardrails.

Following what has been done with the baseline spectral classifier, the same U-Net has been trained with pseudoRGB images which have been extracted as previously described in Section 3.1. Fig. 12 contains the IoU values of the rebuilt images for the 5-class-experiment. As it can be seen, although some classes (Road and Road Marks) are not affected by the reduction of the spectral number of channels, the rest of the classes do suffer a considerable degradation of their metrics.

In ADAS environments, it is of utmost importance to correctly identify traffic signals, cars and pedestrians. As Fig. 13a and b show, our FCN is able to detect those obstacles for the 5-class experiment. However, when we want to determine the class to which those objects belong to, we have to train the FCN using some additional classes (as for Painted Metal in Fig. 13c and d) which, for the mentioned examples, are undersampled. The results obtained are inconclusive and not robust, so we will investigate how to effectively combine spectral classifiers (which can easily cope with unbalanced datasets) with U-Net to improve the segmentation results.

## 5. System prototyping and characterization

As stated above, there are two main processes involved in the implementation of this HSI segmentation system: the preprocessing of the raw images provided by the snapshot camera and the FCN-based AI. In this section we describe in detail how both processing stages have been developed, implemented and tested on three different embedded processing systems: a low-cost single board computer, an embedded GPU SoC and a PSoC with an embedded FPGA. We also give details of the benchmarking setup used for the characterization of all three implementations and we analyze the results obtained.

### 5.1. Benchmark setup

Proper hardware selection is essential for the successful deployment of AI in an edge device. In the particular case of the ADAS/ADS, this decision is conditioned by the need to satisfy strict specifications in terms of processing latency, cost and energy consumption without compromising the overall performance of the algorithm (safety criticality considerations aside). Most current edge-AI deployments rely on SoCs




Fig. 12. IoU values of the rebuilt images for the 5-class-experiment using pseudoRGB images (blue) or HSI (red). (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

including AI-oriented coprocessors (domain specific or even application specific). However, depending on the complexity of the AI model (number of parameters and FLOPS) and the constraints associated with the specific application, the use of cheaper, lower-end devices could provide satisfactory results. Consequently, we have selected three different prototyping platforms to explore the performance of our application: a Raspberry Pi 4B, a NVIDIA Jetson Nano Development Kit and an AMD-Xilinx ZCU104 development board.

The Raspberry Pi 4 Model B (Pi4B) is the leading single-board computer of the Raspberry Pi family with significant enhancement in CPU, GPU and I/O performance. This single-board computer consists of a quad-core Cortex-A72 (ARM v8) 64-bit SoC 1.5 GHz with 8 GB of LPDDR4 SDRAM, which includes a recent upgrade that lets the cores work in turbo-boost mode and reach a peak frequency of 1.8 GHz safely [40]. The NVIDIA Jetson Nano is a small, powerful computer for embedded AI applications with 4 GB 64-bit LPDDR4 SDRAM. The SoC combines a Quad-core ARM Cortex-A57 MPCore processor (1.43 GHz of maximum theoretical frequency) with a Maxwell architecture NVIDIA GPU containing 128 CUDA cores [41]. Finally, the ZCU104 development board includes a Zynq UltraScale+ MPSoC with a quad-core ARM Cortex-A53 processor (1.5 GHz of maximum theoretical frequency) and a dual-core Cortex-R5 real-time processor in the Processing System (PS) connected to a 16 nm FinFET Programmable Logic (PL) with access to a 2 GB 64-bit wide DDR4 external memory [42].

Although all three devices include ARM v8-A architecture Quad-Core Cortex processors (A53, A57 and A72), they do not only differ in CPU frequency but also in other aspects, such as L1 I/D Cache Size, L2 Cache Size, execution order, number of pipeline stages and so on. Cortex A53 is said to be a power-efficient processor when compared to the more powerful A57, while A72, which is the direct successor of the A57, was designed to improve its predecessor in the PPA metric: performance, power and area [43] (see Table 8).

In order to make a fair power and energy consumption comparison, we disabled as many peripherals as possible; this includes, in the case of the Pi4B, onboard LEDs, microHDMI, bluetooth or Wifi, for example. Power measurements have been performed either by software commands or by plugging an external power meter between the power supply and the device (Fig. 14). We have selected a USB-C digital meter from Klein Tools [44] that reliably measures (accuracy of  $\pm 1\%$ ) the voltage and current (with a resolution of 0.01 V and 0.01 A

respectively) that flows through both the Raspberry Pi 4B and the Jetson Nano boards. Furthermore, the *tegrastats* command allows us to extract valuable information about the power consumption of the processing elements (CPU, GPU and overall) of the Jetson Nano, so a more comprehensive analysis can be derived. However, there is a mismatch between the values obtained from the *tegrastats* command and the values read from the digital power meter, an issue that has also been reported previously, such as in [45]. As for the Zynq UltraScale+ MPSoC, power monitoring is carried out with the Infineon IR Power-Center GUI that obtains power rail values through Infineon IRPS5401 power controllers [42]. The onboard Infineon IRPS5401 power controllers are accessed through an I2C connector included in the Infineon USB Cable [46–48].

With regard to the communication infrastructure developed to ease the benchmarking process on multiple devices, in Fig. 14 a router can be seen that is used to create a local area network which allows us to communicate with various prototyping boards via the ssh protocol while their Internet access is granted.

## 5.2. Raw image preprocessing

Recording in real driving conditions implies that the setup of the hyperspectral camera cannot be kept invariable. Moreover, lighting conditions cannot be controlled due to changes in the scenarios, weather conditions and daylight. Additionally, the use of a moving platform (the vehicle) and the presence of moving objects severely constrain the exposition time and, in consequence, the amount of light that can reach the sensor in each situation. The HSI-drive database contains images acquired in very diverse conditions that require different setups of the camera and optics. This diversity makes it more difficult to effectively preprocess the raw images in order to generate hyperspectral cubes with equivalent characteristics that can be used to successfully train AI models. In addition, this preprocessing has to be replicated on-the-fly during the inference phase, so the computational aspects of the preprocessing pipeline have also to be addressed.

The processing stages involved in the pipeline of the raw image preprocessing to obtain the corresponding reflectance hyperspectral cubes depend on the type of application. In our context, the programmed preprocessing algorithm consists of the following four steps: crop and clip, reflectance correction, band extraction, and translation to center (partial demosaicing).

As reported in Section 3, the CMOS sensor has a standard size of  $1088 \times 2048$ . However, as the multispectral filter array (MSFA) does not cover the entire surface of the sensor, the resulting active area has a lower resolution ( $1080 \times 2045$ ) and so the acquired raw images need to be cropped and clipped. After that, the acquired radiance is normalized by taking a white reference frame as the highest response from the sensor, resulting in a reflectance signal which is assumed to be robust under different acquisition conditions. In order to eliminate the static noise of the sensor, bias correction is previously performed by subtracting a dark reference frame from both the image and the white reference frames. The next step, known as band extraction, transforms the 2D representation obtained from the mosaic-shaped MSFA (Fig. 15) into a hyperspectral cube. As a consequence of the snapshot camera technology, each of the 25 pixels of the  $5 \times 5$  mosaic contains spectral information about a different wavelength (Fig. 15) and thus each pixel lacks the spectral content related to the remaining 24 wavelengths. Therefore, to reconstruct from the raw data either a full resolution image ( $1080 \times 2045 \times 25$ ) or a partial resolution image ( $216 \times 409 \times 25$ ) a typical approach is to use an interpolation algorithm in a process known as (partial) demosaicing. In our application, we have opted for a partial demosaicing process which consists of applying linear or bilinear interpolation techniques to interpolate the reflectance value of the missing spectral bands in certain spatial positions which are the center of the mosaic (hence its name, translation to center). Two examples are shown in Fig. 15. On the upper part of the image, we

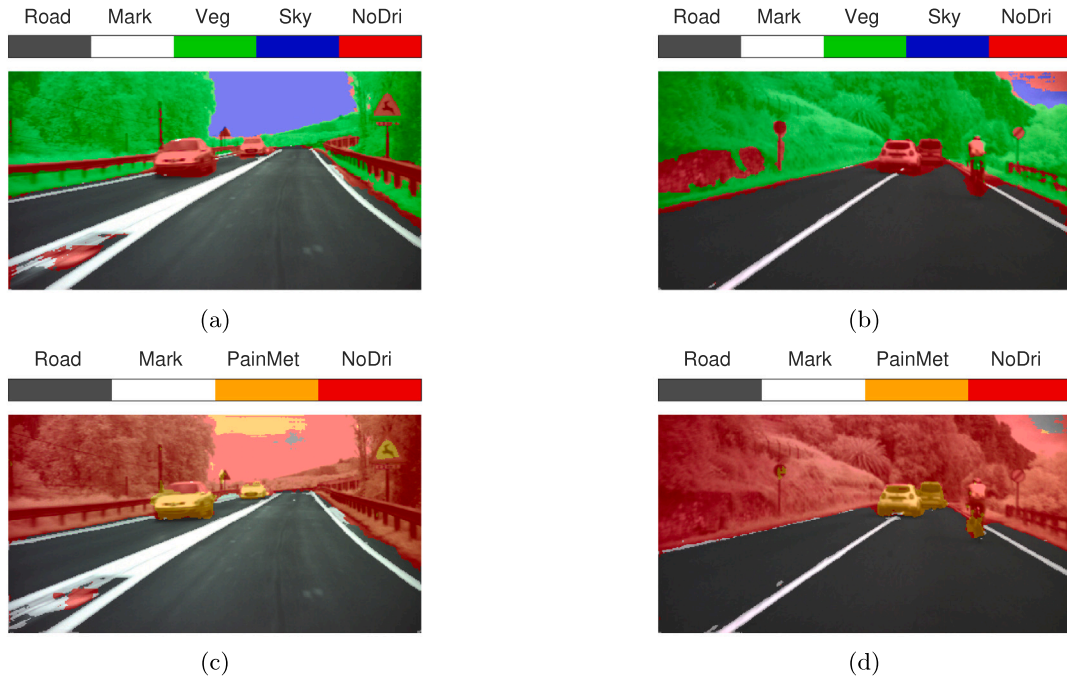


Fig. 13. U-Net segmentation comparison of two selected images: 5-class experiment (a-b) and 3-class experiment + Painted Metal (c-d).




Fig. 14. Power measurement infrastructure and communication with computing boards: 1, Raspberry Pi 4B; 2, Router; 3, Jetson Nano; 4, Power meter and 5, ZCU104.

**Table 8**  
Comparison among the selected platforms focusing on the processors' characteristics.

Properties	Core		
	Cortex-A53	Cortex-A57	Cortex-A72
Revision	v8.0-A	v8.0-A	v8.0-A
Platform	Zynq UltraScale+ MPSoC	Jetson Nano SoM	Raspberry Pi 4B
OS	Petalinux 2022.1 <sup>a</sup>	Ubuntu 18.04.6 LTS	Ubuntu 22.04.1 LTS
Number of cores	4	4	4
Execution	In-order	Out-of-order	Out-of-order
Pipeline depth	8	15+	15+
Superscalar	Yes	Yes	Yes
Neon and FPU	Yes	Yes	Yes
L1 I Cache (per core)	32 KB	48 KiB	48 KiB
L1 D Cache (per core)	32 KB	32 KiB	32 KiB
L2 Cache	1 MB	2 MiB	1 MiB
Clock rate	1.199 GHz	1.479 GHz	1.8 GHz
External DDR4 memory	2 GB	4 GB	8 GB

<sup>a</sup>This OS has been used to work more comfortably with AMD-Xilinx's Tools. However, for TensorFlow Lite model development, we have opted for Ubuntu 20.04.3 LTS.




Fig. 15. Translation to center algorithm.

want to interpolate in the spatial position which corresponds to band 13, the reflectance value that we would obtain if we placed there a band-7 filter. For that purpose, we take into account the values of the surrounding four band-7 pixels, and the distance from the band-13 pixel, and apply a bilinear interpolation to extract that value. We perform this operation to extract all the missing information from the remaining bands but only in the central position of each mosaic. However, as the other example suggests, there are some positions in the contour of the image for which there are not 4 neighboring pixels and so a different treatment (such as linear interpolation) is required.

Differently from what is made in [7], the cube generation contains neither the median filtering nor the band normalization steps. This is because although a prior analysis demonstrated the benefits of using a median filter before training a multi-layer feedforward network, those effects are largely overshadowed in an FCN by the convolutional layers which act as a kind of spatial regularizer and, consequently, the overall latency is greatly reduced. Regarding the normalization, we have performed hardware acceleration on it, including this step as the first layer of U-Net.

5.2.1. Embedded software programming

The raw image preprocessing algorithm has been programmed in C and compiled to be executed as an embedded Linux application in the microprocessors of the benchmark devices as part of the HW/SW co-design for the implementation of the system. To reduce its latency, we have combined thread-level parallelism (OpenMP pragmas) with data-level parallelism (Single instruction multiple data, SIMD, via Neon). OpenMP is an API to develop parallel applications on shared-memory processors, such as embedded systems and accelerators, in a flexible manner [49]. OpenMP grants the programmer control of the thread creation, workload distribution among the threads (not only the scheduling type but also the restriction of a certain task to one thread), thread synchronization and variable attributes (i.e. which variables are shared among threads and which are kept private to avoid data incoherence and race conditions). The way of working with OpenMP is by adding compiler directives or pragmas and modifying environment variables that condition both the compile and runtime behavior of the program in a fork-join model. OpenMP can be combined with Neon (ARM Advanced SIMD architecture) and Floating Point technologies which are fully integrated into the processor and share the processor resources for integer operation, loop control and caching, significantly reducing the area and power cost [50]. SIMD instructions are carried out in the 32 128-bit SIMD/floating-point registers that the AArch64 architecture includes.

To implement those techniques in our software design, on the one hand, some compiler directives have been added to the code after

Table 9

Mean execution time over 1000 iterations of the C image-preprocessing Linux application on the ZCU104 development board combining the use or not of SIMD and MP.

Pipeline step	Use of SIMD-MP			
	00	01	10	11
Image cropping (ms)	6.654	2.883	6.654	2.886
Reflectance correction (ms)	251.803	90.902	50.369	21.519
Band extraction (ms)	53.063	16.217	311.803	13.042
Translation to center (ms)	340.196	162.219	22.586	13.583
Total (ms)	651.716	272.223	111.453	51.030

analyzing how to parallelize the code execution among the different threads. The analysis consisted of identifying parallel regions (coarse-grained parallelism) and then focusing on the parallelizable loops (fine-grained parallelism). In order to parallelize loops, we have studied the dependency between loop indexes or data inside the loops. Besides this, as all threads can modify and access all variables, we have also examined when to change the scope of one of the variables and make them local to each thread to avoid undesirable behavior. Finally, we have checked if the execution time of each iteration varies inside a loop to establish how to distribute the workload among threads. On the other hand, we have also included compiler autovectorization to boost the use of the Neon technology.

To evaluate the extent to which the use of SIMD instructions and multiprocessing accelerates our application, an ablation study has been performed. The results of the four combinations are shown in Table 9.

The combination of SIMD instructions with the thread-level parallelism enabled by the QuadCore ARM-Cortex A we are working with significantly reduces latency with respect to the options where no optimization (12x), or only one of them (5x for MP-only and 2x for SIMD-only approaches), is used. In addition, by analyzing each of the preprocessing stages individually, some conclusions that are consistent with the characteristics of the stages can be drawn.

For the crop and clip stage, the only variable that conditions the execution time is the use or not of multiprocessing. This can be easily explained since this stage only consists of cropping and copying an array where the calculation of the start/end indexes is straightforward.

On the other hand, as previously mentioned, the reflectance correction stage consists of a normalization based on several references, so that a significant acceleration can be achieved with only the use of SIMD instructions, although this is more evident in the median filtering stage, as will be discussed below.

As with the first stage, during partial demosaicing, the image is transformed from a two-dimensional representation to a three-dimensional representation, so the differential factor is whether or not to use multiprocessing. However, since the way in which the indexes for band extraction are calculated requires several mathematical operations that can be parallelized, the use of SIMD instructions additionally reduces the latency.

Of the four stages of preprocessing, band alignment is undoubtedly the one that contains the most mathematical operations, both when choosing the indexes of the pixels selected for interpolation and when performing the interpolation operation itself. This justifies the considerable acceleration observed between using either only SIMD or only opting just for multiprocessing.

5.2.2. Test results

Table 10 shows the mean latency over 1000 iterations of the raw image preprocessing pipeline running on the Cortex A-72 (Raspberry Pi 4B), Cortex A-57 (Jetson Nano), and Cortex A-53 (ZCU104) quad-core processors. According to these figures and if no pipeline schemes are implemented, this processing stage would limit the reachable throughput to 19.10 FPS, 40.56 FPS and 19.84 FPS respectively.

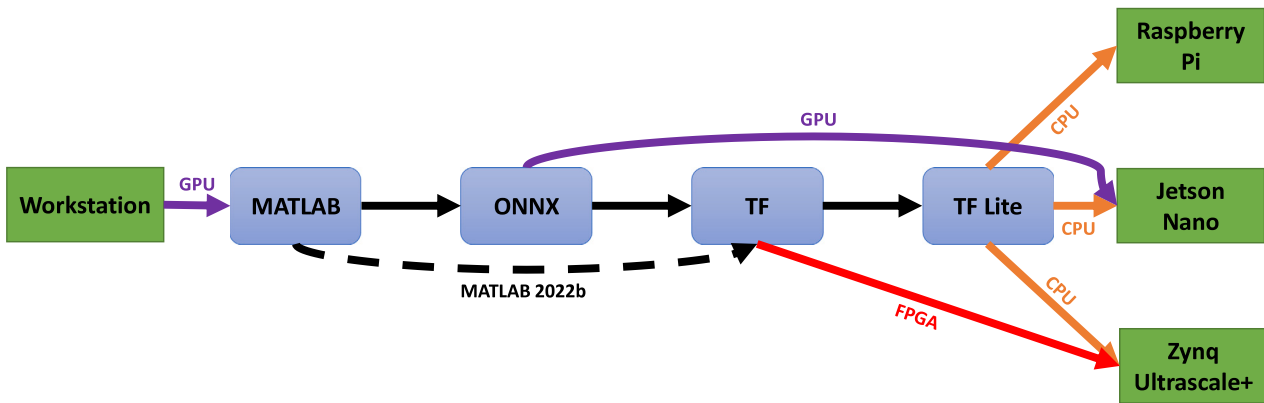


Fig. 16. Path adopted to transfer the FCN from one representation to another.

Table 10

Mean execution time over 1000 iterations of the C image-preprocessing Linux application on the benchmarked devices: Raspberry Pi 4B (RPI), Jetson Nano (JN) and ZCU104 development board (ZCU).

Pipeline step	Device type		
	RPI	JN	ZCU
Image cropping (ms)	6.128	2.309	2.886
Reflectance correction (ms)	13.630	7.729	21.519
Band extraction (ms)	21.124	8.324	13.042
Translation to center (ms)	11.453	6.290	13.583
Total (ms)	<b>52.335</b>	<b>24.652</b>	<b>51.030</b>

### 5.3. FCN inference

The implementation of the FCN has been carried out in several phases, starting with the study of the effect of different quantization levels on the ARM v8 quad-core Cortex processors and ending with the implementation of the optimized models on the AI hardware accelerators, when available, such as the NVIDIA GPU and the AMD-Xilinx DPUs on the FPGA.

#### 5.3.1. Implementation workflow and tools

The design, training and validation of the FCN has been performed with MATLAB's Deep Learning Toolbox. While the most recent version of MATLAB includes tools to directly export deep models, including segmentation FCNs, to other frameworks such as TensorFlow, and even provide tools for the automatic deployment of deep models in FPGAs [51], this was not possible with the Matlab version used in this project (2021b).

Consequently, to link our ML development environment to any implementation workflow, we opted for exporting the MATLAB-generated deep models to the Open Neural Network Exchange (ONNX) representation. ONNX is an open standard format for representing machine learning models that can be then used with a variety of frameworks, tools and compilers. In our workflow, the ONNX representation was exported to Keras/TensorFlow [52] via *onnx2keras* [53]. Fig. 16 shows the flowchart that represents the workflow adopted to transfer the FCN from one representation to another and adapt it to the requirements of both the DL frameworks and the computing platforms under study.

As has been remarked in Section 3.2, we favored the use of a smaller-size FCN architecture over the most accurate one in order to reduce memory footprint and computational load. In an initial approach, we tested the software execution of the FCN directly on the three Cortex microprocessors. With that aim in mind, we used TensorFlow Lite, a library for deploying models on smartphones, microcontrollers and other edge devices. By default, TensorFlow Lite utilizes CPU kernels that are optimized for the Neon instruction subset of ARM, but which are not necessarily optimized for the intensive computational

workload of deep models. Thus, to reduce the final latency as much as possible, we partially used XNNPACK, a highly optimized library of floating-point, neural network inference operators for ARM [54], as the inference engine, which we combined with the multicore execution. Nevertheless, the obtained latency values are not reasonable in ADAS; for this reason, in order to speed up the FCN inference and reduce the memory footprint in all cases, we opted for another strategy and performed post-training quantization using TensorFlow Lite.

As we have prioritized the highest possible speed-up, we have opted for the full integer quantization approach where the whole network (weights, biases, activation functions, input and output) is dynamically quantized (a calibration dataset is therefore needed) using 8-bit integer precision. It is interesting to remark that, although having applied full integer quantization, the size of the data buffer is not reduced fourfold. This is because in the quantization procedure, bias vectors are still quantized as 32-bit data. As is justified in [55], there is a need for the high precision of bias vectors because, as they are added to many output activation functions, otherwise, this could result in an overall bias which could endanger the accuracy of the network. Besides, 8-bit integer products are accumulated using a 32-bit accumulator so 32-bit registers cannot be avoided.

Finally, in order to get the best possible performance out of the benchmark computing platforms, we have implemented the FCN on the dedicated AI coprocessors, i.e. the embedded GPU in the Jetson Nano and the DPUs in the MPSoC's FPGA. In order to make as fair a comparison as possible, the three devices were programmed using their preferred DL framework, i.e. Tensorflow Lite for the Raspberry Pi 4B, Tensor RT for the Jetson Nano, and Vitis AI for the AMD-Xilinx MPSoC. TensorRT, built on the NVIDIA CUDA parallel programming model, is an SDK provided by NVIDIA to perform low latency and high throughput DL inference applications by optimizing trained neural networks and calibrating them for low precision without accuracy degradation. The resultant network can be deployed in diverse platforms, ranging from hyperscale data centers to embedded processors [56]. Vitis AI is a development platform for AI inference on AMD-Xilinx hardware platforms. Vitis AI Quantizer exploits quantization and the VAI Optimizer applies channel pruning techniques to reach the high-throughput and low-latency requirements of ADAS applications. According to [57], by converting the 32-bit floating point weights and activations to an 8-bit integer format, Vitis AI Quantizer can reduce computing complexity without losing prediction accuracy and, as the fixed-point network model requires less memory bandwidth, a faster speed is also provided. After compiling the quantized FCN and the DPU with Vitis AI Compiler, the output product is loaded at runtime in the system composed of the ARM CPU and the DPU accelerator in the MPSoC by a C++ application.

The Deep Learning Processor Unit (DPU) is a configurable computation engine, optimized for convolutional neural networks, which can be customized to some extent by modifying its parallelism degree and logic resource utilization. It is implemented in the programmable logic



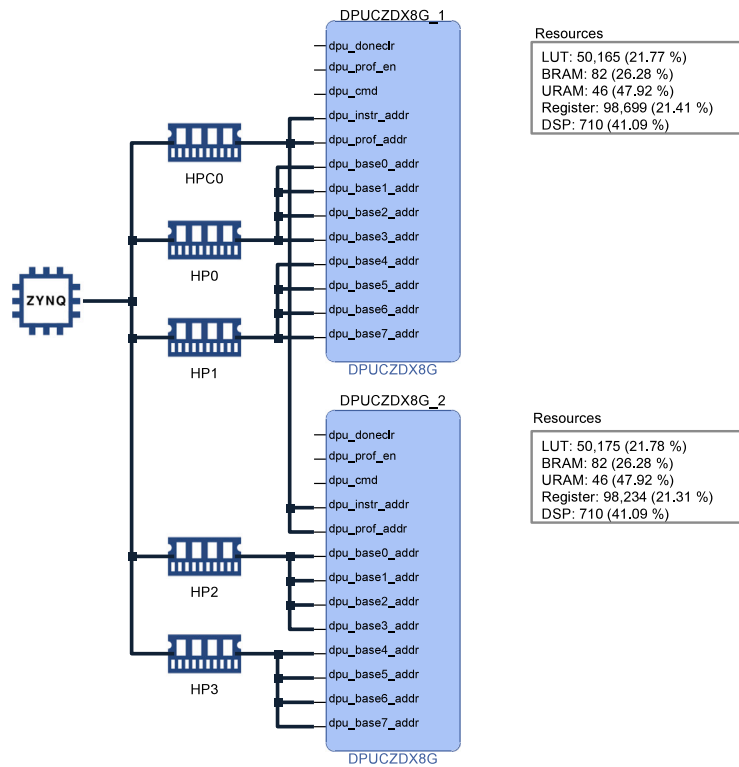


Fig. 17. System diagram of the DPU.

of the MPSoC with direct AXI connections to the processing system. The DPU executes compiled microcode generated from a neural network graph by fetching the instructions from off-chip memory to control how the computing engine behaves. To achieve high throughput, efficiency and reduce external memory bandwidth requirements, on-chip memory is reused as much as possible to buffer input activations, intermediate feature-maps and output activations. The user-configurable parameters are DSP slices, LUT, block RAM and UltraRAM usage. There are even other options for additional functions such as channel augmentation, average pooling, depthwise convolution and softmax. For more detailed information, the reader is referred to [58].

As can be seen in Fig. 17, our implementation contains 2 DPUCZDX8G cores, each of which features a B4096 architecture which means that 4096 operations are performed per DPU clock cycle, which is set to 300 MHz. More specifically, in each clock cycle, 8 pixels of the input feature map (which can have, at most, 16 channels) are multiplied by their respective portions of the 16 convolutional filters, so  $8 \times 16 \times 16$  multiplication operations and sums are performed; this results in a total of 4096 operations. Those three values are architecture-specific parameters known as pixel parallelism, input channel parallelism and output channel parallelism, the latter two sharing the same value. Fig. 17 also shows the logic resources occupied by this implementation.

Although the three DL frameworks allow quantizing models to 8 bits, not all devices are capable of performing computations in this type of representation. The FCN has thus been quantized according to what the respective microarchitectures allow: an 8-bit integer format for the Raspberry Pi, a 16-bit FP format for the Jetson Nano (although as indicated by NVIDIA, TensorRT would choose a higher-precision FP32 kernel if this results in overall lower runtime, or if no low-precision implementation exists), and an 8-bit format representation for the AMD-Xilinx MPSoC.

### 5.3.2. Test results

The implementation of the neural network layers varies slightly from TensorFlow to TensorFlow Lite so, it has first been verified that

there are no significant deviations when exporting the U-Net floating point model from TensorFlow to the Lite version.

The first test has been performed on the software execution of the FCN on the CPU of the three devices both before and after quantization. As can be seen in Table 11, the throughput associated to the FP32 inference latency hardly reaches 1 FPS, while the 8-bit integer implementation produces a notable speed-up, although not enough to meet the application requirements (1.2 FPS to 2.4 FPS).

The second test has been performed on the implementations that use the full power of the SoCs, either with the embedded GPU or the embedded DPU/FPGA. As can be observed in Table 12, the latency is dramatically reduced in both cases, with a x4.47 speed-up in the Jetson Nano and a x22.16 speed-up in the AMD-Xilinx MPSoC. To obtain those values, as each image consists of 18 patches, we have opted for using a 18-patch batch size, although further experiments could be carried out to find the best size to optimize latency-throughput trade-off [59].

Regarding power consumption, the 4.258 W measured in the Raspberry Pi corresponds to the utilization of the whole board although, as previously explained, we have disabled as many peripherals as possible to make a fairer comparison. As for the Jetson Nano, 3.995 W is obtained after applying a correction term (empirically inferred comparing the output of *tegrastats* command and power meter measurement) to the addition of the CPU and GPU consumptions. Properly isolating the chip's consumption from the peripherals allows us to make a finer comparison with the Zynq Ultrascale+ whose 8.8 W also represent the chip's consumption. With respect to energy utilization, the figures show that the higher power consumption of the AMD-Xilinx MPSoC is fully compensated by the latency acceleration achieved, making this the device with the most efficient processing, slightly better than the Jetson Nano.

In a two-stage pipelined processing implementation (image pre-processing + FCN inference), currently achievable throughput limit is imposed by the image preprocessing stage, imposing a limit of around 19,6 FPS. Although this stage could be hardware accelerated in the future to improve throughput, achieved image segmentation latency

**Table 11**

CPU-inference latency characterization of the unquantized and quantized FCN on the three devices: Raspberry Pi 4B (RPI), Jetson Nano (JN) and ZCU104 development board (ZCU).

Framework (Accelerator)	TF Lite (CPU)					
Data quantization	FP32			INT8		
Platform	RPI	JN	ZCU	RPI	JN	ZCU
Measurement						
Latency (s)	1.910	0.835	1.345	0.613	0.416	0.820

**Table 12**

Power, latency and energy consumption measurements regarding the quantization technique (8-bit integer, INT8 or 16/32-bit floating point FP16/FP32), framework (TensorFlow Lite, TF Lite; Tensor RT, TRT and Vitis AI), accelerator and platform (Raspberry Pi 4B (RPI), Jetson Nano (JN) and ZCU104 development board (ZCU)) used.

Framework (Accelerator)	TF Lite (CPU)	TRT (GPU)	Vitis AI (FPGA)
Data quantization	INT8	FP16/FP32	INT8
Platform	RPI	JN	ZCU
Measurement			
Power (W)	4.258	3.995	8.8
Latency (s)	0.613	0.093	0.037
Energy (J)	<b>2.608</b>	<b>0.372</b>	<b>0.326</b>

**Table 13**

U-Net model size variation after applying full integer quantization.

Quantization	3-class experiment				5-class experiment			
	No		Yes		No		Yes	
Buffer type	Memory (KB)	%	Memory (KB)	%	Memory (KB)	%	Memory (KB)	%
Non-data	<b>11.432</b>	8.42	<b>17.108</b>	35.06	11.432	8.41	17.156	35.11
Data	<b>124.392</b>	91.58	<b>31.684</b>	64.94	124.464	91.59	31.708	64.89
Total	135.824	100	48.792	100	135.896	100	48.864	100

(50 ms) is adequate for most ADS/ADAS systems such as forward collision warning systems, emergency breaking systems, trajectory planning systems, etc. [60].

Finally, we present results from the analysis of the memory footprint and the real performance of the segmentation. On Table 13, it is interesting to note how the non-data buffers (responsible for storing, among other things, operators and subgraphs) have the same size for both experiments when quantization has still not been performed on the network. This is due to the fact that the difference between the two models lies solely in the number of filters of the final convolution layer and not in any other operator.

Table 13 also shows that data buffers take up more than 90% of the whole unquantized model. This is what allows the size of the complete model to be reduced considerably when quantization is applied, allowing it to be stored in a device with fewer memory resources. However, it has to be mentioned that there is a 50% increase in the size of the non-data buffers after quantization. This fact is a consequence of the inclusion of new operators (scaling or casting, for example) which are needed for the aforementioned process. As for the segmentation variations, we show in Table 14 the recall, precision and IoU after the different quantization approaches.

Table 14 includes some interesting results. With regard to the 3-class experiment, the 8-bit integer Raspberry Pi 4B and the AMD-Xilinx MPSoC quantization approaches (each device with its own quantization technique) offer similar and even better performance in several cases than the 16/32-bit floating point Jetson Nano. Furthermore, these results also verify that the use of overlapping patches improves the segmentation recall/precision. Nevertheless, the results from the 5-class experiment show that having added two classes to the initial segmentation application increases the challenge difficulty to such an extent that the best results are now obtained with the Jetson Nano. These results highlight the need to either explore alternative quantization or finetuning techniques or directly train the neural network in such a way that it is aware that it will be later quantized, a process known as Quantization Aware Training (QAT) [57].

Although metrics such as those listed in Table 14 are a good indicator of the quality of inference, it is essential to inspect the inference results in the images, especially when dealing with weakly labeled datasets such as the one we are working with. Otherwise, the prediction of the neural network in the unlabeled pixels, which are equally as important as the other ones, is unknown.

For the visual inspection of the results, we have chosen the same images as before. Although, at first sight, the three images all look similar, we will thoroughly analyze them one by one. As for the urban image, the output from the AMD-Xilinx MPSoC (Fig. 18(i)) is that which has the smallest white spot in the foreground; however, it is also the one that fails to perceive the left side thin road lines. The output from Raspberry Pi 4B and Jetson Nano are almost identical (Figs. 18(c) and 18(f)). Regarding the road landscape, the segmentation of the Raspberry Pi 4B (Fig. 18(b)) is the one that makes the smallest error, followed by the Jetson Nano's (Fig. 18(e)) in the vegetation to the right of the road sign and in the sky/vegetation intersections. However, they have less ability than the AMD-Xilinx MPSoC (Fig. 18(h)) to detect objects located at medium-long distances; this is evident in the sharpness of both the signal on the left and the guardrail underneath it, in the couple who is going for a walk and in the house at the back. Finally, in relation to the highway scene, which is challenging, once again the Raspberry Pi 4B and the Jetson Nano (Figs. 18(c), 18(f)) show a different output compared to the AMD-Xilinx MPSoC (Fig. 18(i)): while the former better identify the lane lights before the tunnel entry, the latter is the only one to detect the road on the right side and to partially predict the road on the left. Nevertheless, the overall segmentation is extremely good.

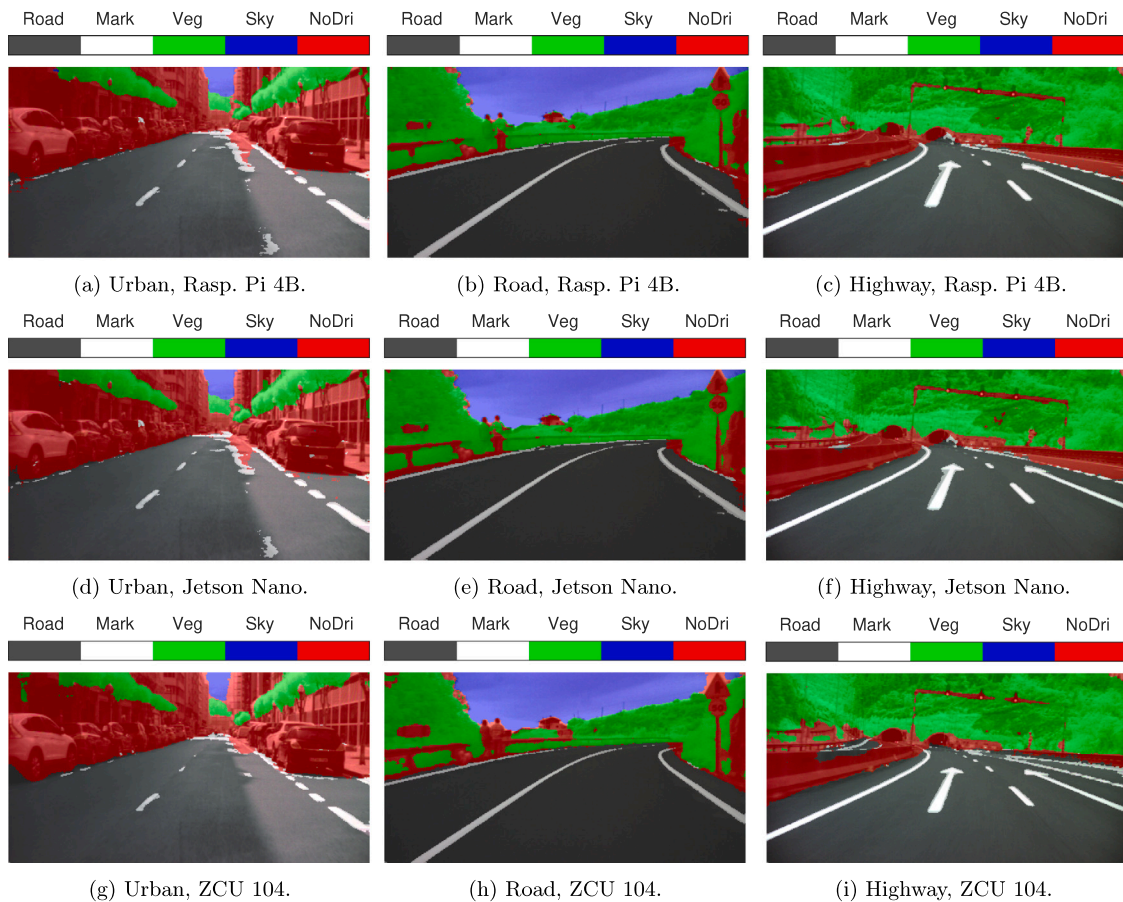
## 6. Conclusions

HSI contains relevant spectral information that can help to develop more capable intelligent vision systems. In the field of ADAS and ADS, HSI can provide sufficient information to produce the correct segmentation of relevant items in driving scenarios and thus provide

**Table 14**

Recall, accuracy and IoU (Intersection over Union) of the quantized neural network inference running on the three tested computing platforms: Raspberry Pi 4B (RPI), Jetson Nano (JN) and ZCU104 development board (ZCU).

Rebuilt									
Metric	Recall			Precision			IoU		
Class name	Device								
	RPI	Nano	ZCU	RPI	Nano	ZCU	RPI	Nano	ZCU
Road	98.28	98.19	97.97	95.44	95.72	97.11	93.87	94.06	95.19
Road marks	86.13	84.42	86.71	80.98	79.91	81.12	71.63	69.64	72.15
Non-drivable	92.87	93.45	95.59	97.79	97.71	97.40	90.96	91.44	93.21
Overall	95.86	95.98	96.74	95.84	95.96	96.77	92.27	92.50	93.78
Mean	92.42	92.02	93.42	91.40	91.11	91.88	85.49	85.05	86.85
Weighted	87.08	85.61	87.78	82.75	81.80	82.88	73.89	72.16	74.52
Road	90.22	93.26	90.37	99.24	98.75	99.05	89.60	92.17	89.59
Road marks	78.66	75.91	78.30	78.53	78.61	78.69	64.74	62.92	64.60
Vegetation	96.90	95.86	96.02	91.85	95.73	95.63	89.23	91.93	91.99
Sky	97.64	97.51	95.40	92.85	93.56	95.40	90.81	91.38	91.21
Other	78.47	84.01	85.50	57.29	64.44	57.34	49.51	57.40	52.25
Overall	90.58	92.62	91.08	92.34	93.61	93.20	84.41	87.42	85.66
Mean	88.38	89.31	89.12	83.95	86.22	85.22	76.78	79.16	77.93
Weighted	86.92	86.25	85.67	82.65	83.96	82.28	74.58	75.06	73.06



**Fig. 18.** Results from the 5-class quantized U-Net inference on the benchmarked devices of three diverse images.

the vehicles with meaningful data for scene understanding. However, ADAS and ADS require capturing and processing HSI at video rates, which poses an additional challenge to the task of image segmentation. Nowadays, hyperspectral imaging at video rates in small formats is possible mainly by using snapshot cameras that incorporate standard CMOS imaging sensors with filter-on-chip technology. Filter-array technology constrains, to some extent, the quality of the captured spectral and spatial information, especially in outdoor uncontrolled environments, as is the case in real driving situations. Consequently, to achieve

good classification performance and robustness, it is necessary to apply additional techniques beyond the mere use of raw spectral information.

It is known that the use of spatial information via convolution operations allows for obtaining a more robust segmentation of images by applying deep-learning techniques, but it is at the cost of a huge increase in the complexity of the models. However, the requirements of ADAS applications impose tight constraints on the complexity of the AI models, since they must be deployed in compliance with the low cost, low latency and low power requirements of embedded computing

platforms. In this regard, we demonstrate that satisfactory segmentation results can be obtained for low/medium-complexity use cases when combining spectral information from a snapshot camera with the spatial information obtained with light encoder–decoder FCN models. In particular, we provide evidence that with a careful design and a hyperparameter optimization process, it is possible to achieve good performance with FCNs of sizes in the range of tens of thousands of parameters.

We provide a benchmark of designed segmentation systems by the deployment on three significant embedded computing platforms: a multicore microprocessor, a SoC featuring an embedded GPU, and a PSoC containing an FPGA. We show that the FPGA-based PSoC approach is superior in terms of both latency and energy efficiency. The PSoC implementation allows for achieving almost 20 fps of processing throughput with a simple two-stage processing pipeline (preprocessing stage and inference stage), which meets the latency requirements of the ADAS industry standards. In fact, it is important to highlight that, unlike other works published in this field of application, our tested prototypes incorporate not only the acceleration of the neural network inference via the AI dedicated coprocessor, but also the necessary preprocessing of the acquired raw HSI which, in fact, reveals itself to be a bottleneck that needs to be taken into account in the image segmentation pipeline. In this sense, we also show that accelerating the preprocessing stage by taking advantage of both the thread-level and the data-level parallelism of modern microprocessor architectures in the writing of the code, makes the difference in terms of throughput.

Although the presented results are promising, HSI-based intelligent vision research is still at its early stages and there is much room for improvement. Future work will focus on extending the capabilities of the system at both the algorithmic and the implementation levels. It is expected that the ML-based algorithmic development and model training will be improved as a result of the availability of a greater amount of data. In this sense, the recently published new v2.0 release of the HSI-Drive dataset, which provides 496 additional labeled images extracted from new recordings performed in Winter and Autumn, will facilitate the training of models thanks to the increase of data available, particularly for the minority classes and for challenging scenarios such as rainy weather, low-lighting environments etc. In this regard, it will be necessary to investigate how to extract more meaningful spectral features that would be able to enhance class separability and how to effectively combine this information with the spatial features at different scales.

Regarding the deployment on embedded computing platforms that fit ADAS/ADS application requirements, on the one hand, it is necessary to further explore optimization techniques regarding model quantization, pruning, the use of separable convolutions etc. On the other hand, we are assessing the possibility of designing a more customized AI coprocessor to further accelerate the data flow and reduce resource occupation in order to enhance the system throughput and reduce power consumption.

### Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

### Data availability

The HSI data will be made available on request. The whole dataset can be downloaded from <https://ipaccess.ehu.es/HSI-Drive/>. The code will be uploaded to <https://gitlab.com/EHU-GDED/FLOSS/HSI-Drive>.

### References

- [1] M. Govender, K. Chetty, H. Bulcock, A review of hyperspectral remote sensing and its application in vegetation and water resource studies, *Water Sa* 33 (2) (2007) 145–151.
- [2] M. Colomb, P. Duthon, F. Bernardin, Spectral reflectance characterization of the road environment to optimize the choice of autonomous vehicle sensors, in: 2019 IEEE Intelligent Transportation Systems Conference, ITSC, IEEE, 2019, pp. 1085–1090.
- [3] K. Weikl, D. Schroeder, W. Stechele, Potentials of combined visible light and near infrared imaging for driving automation, in: *Electronic Imaging Conference*, 2022.
- [4] Y. Huang, Q. Shen, Y. Fu, S. You, Weakly-supervised semantic segmentation in cityscape via hyperspectral image, in: *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2021, pp. 1117–1126.
- [5] N. Pinchon, O. Cassagnol, A. Nicolas, F. Bernardin, P. Leduc, J.-P. Tarel, R. Brémond, E. Bercier, J. Brunet, All-weather vision for automotive safety: which spectral band? in: *International Forum on Advanced Microsystems for Automotive Applications*, Springer, 2018, pp. 3–15.
- [6] C. Winkens, F. Sattler, D. Paulus, Deep dimension reduction for spatial-spectral road scene classification, *Electron. Imaging* (2019) 49.1–49.9.
- [7] J. Gutiérrez-Zaballa, K. Basterretxea, J. Echanobe, M.V. Martínez, I. del Campo, Exploring fully convolutional networks for the segmentation of hyperspectral imaging applied to advanced driver assistance systems, in: *Design and Architecture for Signal and Image Processing: 15th International Workshop, DASIP 2022*, Budapest, Hungary, June 20–22, 2022, *Proceedings*, Springer-Verlag, Berlin, Heidelberg, 2022, pp. 136–148, [http://dx.doi.org/10.1007/978-3-031-12748-9\\_11](http://dx.doi.org/10.1007/978-3-031-12748-9_11).
- [8] K. Basterretxea, V. Martínez, J. Echanobe, J. Gutiérrez-Zaballa, I. Del Campo, HSI-drive: A dataset for the research of hyperspectral image processing applied to autonomous driving systems, in: 2021 IEEE Intelligent Vehicles Symposium, IV, 2021, pp. 866–873, <http://dx.doi.org/10.1109/IV48863.2021.9575298>.
- [9] H. Alhajja, S. Mustikovela, L. Mescheder, A. Geiger, C. Rother, Augmented reality meets computer vision: Efficient data generation for urban driving scenes, *Int. J. Comput. Vis. (IJCV)* (2018).
- [10] G.J. Brostow, J. Shotton, J. Fauqueur, R. Cipolla, Segmentation and recognition using structure from motion point clouds, in: *ECCV* (1), 2008, pp. 44–57.
- [11] G.J. Brostow, J. Fauqueur, R. Cipolla, Semantic object classes in video: A high-definition ground truth database, *Pattern Recognit. Lett.* xx (x) (2008) xx.
- [12] M. Cordts, M. Omran, S. Ramos, T. Rehfeld, M. Enzweiler, R. Benenson, U. Franke, S. Roth, B. Schiele, The cityscapes dataset for semantic urban scene understanding, in: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2016, pp. 3213–3223.
- [13] B. Kim, J. Yim, J. Kim, Highway driving dataset for semantic video segmentation, 2020, arXiv preprint [arXiv:2011.00674](https://arxiv.org/abs/2011.00674).
- [14] L.-C. Chen, G. Papandreou, I. Kokkinos, K. Murphy, A.L. Yuille, Deeplab: Semantic image segmentation with deep convolutional nets, atrous convolution, and fully connected crfs, *IEEE Trans. Pattern Anal. Mach. Intell.* 40 (4) (2017) 834–848.
- [15] M. Orsic, I. Kreso, P. Bevandic, S. Segvic, In defense of pre-trained imagenet architectures for real-time semantic segmentation of road-driving images, in: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2019, pp. 12607–12616.
- [16] M. Fan, S. Lai, J. Huang, X. Wei, Z. Chai, J. Luo, X. Wei, Rethinking BiSeNet for real-time semantic segmentation, in: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2021, pp. 9716–9725.
- [17] P. Li, X. Dong, X. Yu, Y. Yang, When humans meet machines: Towards efficient segmentation networks, in: *The 31st British Machine Vision Virtual Conference*, 2020.
- [18] E. Courdier, F. Fleuret, Real-time segmentation networks should be latency aware, in: *Proceedings of the Asian Conference on Computer Vision*, 2020.
- [19] G.A. Fricker, J.D. Ventura, J.A. Wolf, M.P. North, F.W. Davis, J. Franklin, A convolutional neural network classifier identifies tree species in mixed-conifer forest from hyperspectral imagery, *Remote Sens.* 11 (19) (2019) 2326.
- [20] M. Taghizadeh, A.A. Gowen, C.P. O'Donnell, Comparison of hyperspectral imaging with conventional RGB imaging for quality evaluation of *Agaricus bisporus* mushrooms, *Biosyst. Eng.* 108 (2) (2011) 191–194.
- [21] S. Seidlitz, J. Sellner, J. Odenthal, B. Özdemir, A. Studier-Fischer, S. Knödler, L. Ayala, T.J. Adler, H.G. Kenngott, M. Tizabi, et al., Robust deep learning-based semantic organ segmentation in hyperspectral images, *Med. Image Anal.* (2022) 102488.
- [22] P. Borges, T. Peynot, S. Liang, B. Arain, M. Wildie, M. Minareci, S. Lichman, G. Samvedi, I. Sa, N. Hudson, et al., A survey on terrain traversability analysis for autonomous ground vehicles: Methods, sensors, and challenges, *Field Robot.* 2 (1) (2022) 1567–1627.
- [23] J. Lu, H. Liu, Y. Yao, S. Tao, Z. Tang, J. Lu, Hsi road: A hyper spectral image dataset for road segmentation, in: 2020 IEEE International Conference on Multimedia and Expo, ICME, 2020, pp. 1–6, <http://dx.doi.org/10.1109/ICME46284.2020.9102890>.

- [24] D.C. Liyanage, R. Hudjakov, M. Tamre, Hyperspectral imaging methods improve RGB image semantic segmentation of unstructured terrains, in: 2020 International Conference Mechatronics Systems and Materials, MSM, 2020, pp. 1–5, <http://dx.doi.org/10.1109/MSM49833.2020.9201738>.
- [25] J. Herweg, J. Kerekes, M. Eismann, Separability between pedestrians in hyperspectral imagery, *Appl. Opt.* 52 (6) (2013) 1330–1338, <http://dx.doi.org/10.1364/AO.52.001330>, <https://opg.optica.org/ao/abstract.cfm?URI=ao-52-6-1330>.
- [26] L. Cavigelli, D. Bernath, M. Magno, L. Benini, Computationally efficient target classification in multispectral image data with deep neural networks, in: Target and Background Signatures II, Vol. 9997, SPIE, 2016, pp. 191–202.
- [27] C. Winkens, V. Adams, D. Paulus, Automatic shadow detection using hyperspectral data for terrain classification, *Electron. Imaging* (2019) 31–31.6.
- [28] C. Winkens, V. Kobelt, D. Paulus, Robust features for snapshot hyperspectral terrain classification, in: 17th Int. Conf. Computer Analysis of Images and Patterns, 2017, pp. 16–27.
- [29] C. Winkens, D. Paulus, Context aware hyperspectral scene analysis, in: Int. Symp. Electronic Imaging Science and Technology, 2018, pp. 346.1–346.7.
- [30] Photonfocus: MV1-D2048x1088-HS02-96-G2. <https://www.photonfocus.com/products/camerafinder/camera/mv1-d2048x1088-hs02-96-g2>.
- [31] S. Wang, W. Chen, S.M. Xie, G. Azzari, D.B. Lobell, Weakly supervised deep learning for segmentation of remote sensing imagery, *Remote Sens.* 12 (2) (2020) 207.
- [32] G. Forestier, J. Inglada, C. Wemmert, P. Gançarski, Comparison of optical sensors discrimination ability using spectral libraries, *Int. J. Remote Sens.* 34 (7) (2013) 2327–2349.
- [33] J. Bai, W. Shi, Z. Xiao, T.A.A. Ali, F. Ye, L. Jiao, Achieving better category separability for hyperspectral image classification: A spatial-spectral approach, *IEEE Trans. Neural Netw. Learn. Syst.* (2023).
- [34] Q. Du, H. Yang, Similarity-based unsupervised band selection for hyperspectral image analysis, *IEEE Geosci. Remote Sens. Lett.* 5 (4) (2008) 564–568.
- [35] L. Van der Maaten, G. Hinton, Visualizing data using t-SNE, *J. Mach. Learn. Res.* 9 (11) (2008).
- [36] O. Ronneberger, P. Fischer, T. Brox, U-Net: Convolutional networks for biomedical image segmentation, in: International Conference on Medical Image Computing and Computer-Assisted Intervention, Springer, 2015, pp. 234–241.
- [37] G.-J. Son, D.-H. Kwak, M.-K. Park, Y.-D. Kim, H.-C. Jung, U-net-based foreign object detection method using effective image acquisition system: A case of almond and green onion flake food process, *Sustainability* 13 (24) (2021) 13834.
- [38] X. Cui, K. Zheng, L. Gao, B. Zhang, D. Yang, J. Ren, Multiscale spatial-spectral convolutional network with image-based framework for hyperspectral imagery classification, *Remote Sens.* 11 (19) (2019) 2220.
- [39] H. Iqbal, Harisqbal88/plotneuralnet v1. 0.0, 2018, URL: <https://doi.org/10.5281/Zenodo>.
- [40] R. Pi, Raspberry pi 4 model b datasheet, 2019, <https://datasheets.raspberrypi.com/rpi4/raspberry-pi-4-datasheet.pdf>.
- [41] NVIDIA, NVIDIA Jetson Nano system-on-module, 2022, [https://developer.download.nvidia.com/assets/embedded/secure/jetson/Nano/docs/JetsonNano\\_DataSheet\\_DS09366001v1.1.pdf?MpmY78SZFEFH-jJG2aGI3bYEW75LZBVZJnkH7Bx-4huO6qS6b4u-9T330-jPr9XU3MLlBnF\\_lfY5GTVCIn7wUCwW5pHzJouudAa9FcSOHi1CB0qyQ40PzYgABdhZ9d-smLoLnUV67\\_e8KvMzRjSMahO2ebQ8RZOxboLa\\_GzcnQYaMtBGVa94dW67a7g&t=eyJscyl6ImdzW8iLCJsc2QioiJodHRwczpL1wvd3d3Lmdvb2dsZS5j21cYJ9](https://developer.download.nvidia.com/assets/embedded/secure/jetson/Nano/docs/JetsonNano_DataSheet_DS09366001v1.1.pdf?MpmY78SZFEFH-jJG2aGI3bYEW75LZBVZJnkH7Bx-4huO6qS6b4u-9T330-jPr9XU3MLlBnF_lfY5GTVCIn7wUCwW5pHzJouudAa9FcSOHi1CB0qyQ40PzYgABdhZ9d-smLoLnUV67_e8KvMzRjSMahO2ebQ8RZOxboLa_GzcnQYaMtBGVa94dW67a7g&t=eyJscyl6ImdzW8iLCJsc2QioiJodHRwczpL1wvd3d3Lmdvb2dsZS5j21cYJ9).
- [42] Xilinx, ZCU104 evaluation board, 2018, [https://www.xilinx.com/support/documents/boards\\_and\\_kits/zcu104/ug1267-zcu104-eval-bd.pdf](https://www.xilinx.com/support/documents/boards_and_kits/zcu104/ug1267-zcu104-eval-bd.pdf).
- [43] AnandTech, ARM reveals cortex-a72 architecture details, 2015, <https://www.anandtech.com/show/9184/arm-reveals-cortex-a72-architecture-details>, 10.10.2022.
- [44] K. Tools, USB digital meter, USB-A and USB-C, 2021, <https://www.kleintools.com/catalog/usb-digital-meters/usb-digital-meter-usb-and-usb-c>.
- [45] S. Holly, A. Wendt, M. Lechner, Profiling energy consumption of deep neural networks on nvidia jetson nano, in: 2020 11th International Green and Sustainable Computing Workshops, IGSC, IEEE, 2020, pp. 1–6.
- [46] Infineon, IRPS5401 PMIC flexible power management unit, 2021, [https://www.infineon.com/dgdl/Infineon-IRPS5401M-DataSheet-v02\\_06-EN.pdf?fileId=5546d4625cc9456a015cd69d402139db](https://www.infineon.com/dgdl/Infineon-IRPS5401M-DataSheet-v02_06-EN.pdf?fileId=5546d4625cc9456a015cd69d402139db).
- [47] Infineon, USB005 user guide, 2014, <https://www.infineon.com/cms/en/product/power/dc-dc-converters/digital-multiphase-controllers/gang-programmers/usb005/>.
- [48] Xilinx, HW-Z1-ZCU104 evaluation board (XCZU7EV-2FFVC1156) schematic, 2018, <https://www.xilinx.com/products/boards-and-kits/zcu104.html#resources>.
- [49] OpenMP, About us, 2022, <https://www.openmp.org/about/about-us/>.
- [50] ARM, Learn the architecture - Introducing Neon, 2020, <https://developer.arm.com/documentation/102474/0100/?lang=en>.
- [51] MATLAB, Deep learning toolbox, 2022, <https://es.mathworks.com/products/deep-learning.html>.
- [52] M. Abadi, A. Agarwal, P. Barham, E. Brevdo, Z. Chen, C. Citro, G.S. Corrado, A. Davis, J. Dean, M. Devin, S. Ghemawat, I. Goodfellow, A. Harp, G. Irving, M. Isard, Y. Jia, R. Jozefowicz, L. Kaiser, M. Kudlur, J. Levenberg, D. Mané, R. Monga, S. Moore, D. Murray, C. Olah, M. Schuster, J. Shlens, B. Steiner, I. Sutskever, K. Talwar, P. Tucker, V. Vanhoucke, V. Vasudevan, F. Viégas, O. Vinyals, P. Warden, M. Wattenberg, M. Wicke, Y. Yu, X. Zheng, TensorFlow: Large-scale machine learning on heterogeneous systems, 2015, <https://www.tensorflow.org/>, Software available from tensorflow.org.
- [53] G. Malivenko, Onnx2keras 0.0.24, 2021, <https://pypi.org/project/onnx2keras/>.
- [54] Tensorflow, XNNPACK backend for TensorFlow Lite, 2013, <https://github.com/tensorflow/tensorflow/tree/master/tensorflow/lite/delegates/xnnpack>.
- [55] B. Jacob, S. Kligys, B. Chen, M. Zhu, M. Tang, A. Howard, H. Adam, D. Kalenichenko, Quantization and training of neural networks for efficient integer-arithmetic-only inference, in: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, 2018, pp. 2704–2713.
- [56] NVIDIA, NVIDIA TensorRT, 2022, <https://developer.nvidia.com/tensorrt>.
- [57] Xilinx, Quantizing the model, 2021, [https://www.xilinx.com/html\\_docs/vitis\\_ai/1\\_4/quantize.html#uim1570695919827](https://www.xilinx.com/html_docs/vitis_ai/1_4/quantize.html#uim1570695919827).
- [58] Xilinx, DPUCZDX8G for zynq UltraScale+ MPSoCs product guide (PG338), 2022, <https://docs.xilinx.com/r/en-US/pg338-dpu?tocId=Bd4R4bhnWgMYE6wUISXDLw>.
- [59] J. Hanhirona, T. Kämäräinen, S. Seppälä, M. Siekkinen, V. Hirvisalo, A. Ylä-Jääski, Latency and throughput characterization of convolutional neural networks for mobile computer vision, in: Proceedings of the 9th ACM Multimedia Systems Conference, 2018, pp. 204–215.
- [60] E. NCAP, European new car assessment programme (Euro NCAP), 2021, <https://cdn.euroncap.com/media/62794/euro-ncap-aeb-c2c-test-protocol-v303.pdf>, 03.04.2023.



**Jon Gutiérrez-Zaballa** received the Bachelor's degree in Physics and Electronic Engineering from the University of the Basque Country in 2019 and the M.Sc. on Advanced Electronic Systems at the same university in 2021. He is currently working as a Ph.D. candidate in Physical Engineering at the Department of Electronics Technology of the University of the Basque Country. His current research interests mainly concern image processing, computational intelligence (data mining, Artificial Neural Networks, Convolutional Neural Networks) and software-hardware embedded co-design on PSoCs for Advanced Driver Assistance Systems.



**Koldo Basterretxea** received the Licenciado degree in physics (M.Sc.) with specialization in electronics and automatic control and the Ph.D. degree in physics from the University of the Basque Country (UPV/EHU), Spain, in 1994 and 2002, respectively. Currently, he is with the Department of Electronics Technology, Faculty of Engineering of Bilbao, UPV/EHU, where he has been an Associate Professor since 2009 and coordinates the Master's Program in Advanced Electronic Systems. Dr. Basterretxea has been a visiting researcher at the University of Edinburgh School of Engineering in 2009 and 2012. His research interests include design of adaptive systems on FPGAs, hardware acceleration of numerical matrix algebra and optimization algorithms, and design of application and domain specific processors for high-speed embedded control, machine learning and computational intelligence applications.



**Javier Echanobe** (M'06) received the Licenciado degree in physics from the University of the Basque Country (UPV/EHU), Spain, and the Ph.D. degree from the University of Navarra, Pamplona, Spain, in 1990 and 1998, respectively. He was a Predoctoral Researcher (granted by the Basque Government) from 1992 to 1996. He has been an Associate Professor in the Department of Electricity and Electronics, UPV/EHU, from 1999 to 2009. Since 2009 he is a Senior Lecturer in that department. His research interests focus on (1) digital electronics: embedded systems, reconfigurable FPGAs, DSPs, SoPC; (2) computational intelligence: artificial neural networks, fuzzy systems, and genetic algorithms; (3) ubiquitous computing: ambient intelligence, intelligent environments, intelligent vehicles. Dr. Echanobe has published many papers in international journals and conferences in most of those areas.




**Maria Victoria Martínez** received the Licenciado degree in physics with specialization in electronics and automatics in 1987, and the Ph.D. degree in physics in 2002, both from the University of the Basque Country (UPV/EHU), Bilbao, Spain. She was with the Department of Electricity and Electronics of the UPV/EHU from 1987 to 1991, first as a Predoctoral Researcher and then as a Lecturer. She worked as an informatics analyst from 1991 to 1995, participating in the development of communications projects. Currently, she is a Lecturer at the Department of Electricity and Electronics of the UPV/EHU. Her research interests mainly concern machine learning, data mining, artificial neural networks, computational intelligence and intelligent environments.



**Unai Martínez-Corral** received the Ingeniería Técnica Industrial B.Sc. in 2013, and the Advanced Electronic Systems M.Sc. in 2015, both from the University of the Basque Country (UPV/EHU), Bilbao, Basque Country, Spain. Since 2016, he is a Ph.D. candidate in the Engineering Physics programme with the Digital Electronics Design Group at the UPV/EHU. His research interests include open source tooling for the design, verification and implementation of adaptive systems on FPGA, hardware design for high-speed real-time controllers, and hardware–software co-design. In Q1 2021, he was awarded a Google Open Source Peer Bonus for his contributions to EDA tooling projects.



**Óscar Mata-Carballera** received the Bachelor's degree in Industrial Electronics and Automatics Engineering from the University of A Coruña in 2016, the M.Sc. on Advanced Electronic Systems from the University of the Basque Country in 2018 and the Ph.D. in Physical Engineering at the same university under the University Teachers Training Program from the Government of Spain (FPU) in 2022 with honors. He is currently working as a Ph.D. researcher at the Department of Electricity and Electronics of the University of the Basque Country. His current research interests mainly cover computational intelligence algorithms (mainly shallow and deep Extreme Learning Machines, Fuzzy Logic and Convolutional Neural Networks) as well as image processing, hardware–software partitioning, hardware design on VHDL, FPGAs and advanced driving assistance systems.



**Inés del Campo** was born in Buenos Aires, Argentina, in 1961. She received the Licenciado degree in physics with specialization in electronics and automatics in 1987 and the Ph.D. degree in physics, in 1993, both from the University of the Basque Country (UPV/EHU), Bilbao, Spain. Currently she is a Full Professor in the Electricity and Electronics Department of the Faculty of Sciences and Technology of the UPV/EHU. She has published articles in international journals and conferences in the areas of electronics, computational intelligence, intelligent control, and intelligent transportation systems, among others. Her research interests mainly concern system-on-chip (SOC) design, reconfigurable hardware, machine learning, artificial neural networks, deep learning, and data mining. She is also interested in the design of advanced driver assistance systems (ADAS), ride comfort, autonomous driving, and hyperspectral imaging for scene understanding.