

# DEEP LEARNING FOR SEMANTIC SEGMENTATION OF REMOTE SENSING IMAGES WITH RICH SPECTRAL CONTENT

A. Ben Hamida<sup>1,2</sup>, A. Benoit<sup>1</sup>, P. Lambert<sup>1</sup>, L. Klein<sup>1</sup>, C. Ben Amar<sup>2</sup>, N. Audebert<sup>3,4</sup>, S. Lefèvre<sup>4</sup>

<sup>1</sup> Univ. Savoie Mont Blanc, LISTIC, F-74000 Annecy, France  
(amina.ben-hamida,alexandre.benoit,patrick.lambert,louis.klein)@univ-smb.fr

<sup>2</sup> REGIM, Ecole Nationale d'Ingénieurs de Sfax, 3038 Sfax, Tunisie  
chokri.benamar@ieee.org

<sup>3</sup> ONERA, *The French Aerospace Lab*, 91761 Palaiseau, France

<sup>4</sup> Univ. Bretagne-Sud, UMR 6074, IRISA, 56000 Vannes, France  
(nicolas.audebert,sebastien.lefevre)@irisa.fr

## ABSTRACT

With the rapid development of Remote Sensing acquisition techniques, there is a need to scale and improve processing tools to cope with the observed increase of both data volume and richness. Among popular techniques in remote sensing, Deep Learning gains increasing interest but depends on the quality of the training data. Therefore, this paper presents recent Deep Learning approaches for fine or coarse land cover semantic segmentation estimation. Various 2D architectures are tested and a new 3D model is introduced in order to jointly process the spatial and spectral dimensions of the data. Such a set of networks enables the comparison of the different spectral fusion schemes. Besides, we also assess the use of a “noisy ground truth” (i.e. outdated and low spatial resolution labels) for training and testing the networks.

**Index Terms**— Remote Sensing, Multispectral, Deep Learning, Semantic Segmentation, Noisy Training.

## 1. INTRODUCTION

Automated land cover mapping based on satellite image analysis and classification is a well-known challenge which is of a great interest for many fields, such as agriculture [1] and risk monitoring [2]. The recently launched Sentinel-2 satellite constellation provides a richer content in spatial, spectral and temporal domains, and produces a huge amount of images to process daily. In this context, Deep Learning (DL) appears as an appealing alternative to traditional *shallow* classification approaches to deal with such a massive amount of data. A DL architecture is a deep artificial neural network composed of a hierarchical succession of neuron layers performing linear and non-linear processing. Mimicking the human brain behavior, the network tuning (typically millions of parameters) is automatically performed thanks to a supervised training process on large datasets that are generally associated to

some “ground truth” knowledge. However, the ground truth quality is essential to reach satisfying performances.

Some recent works already use deep neural nets to process remotely sensed images. In [3], the authors compared different well-established deep architectures (AlexNet, AlexNet-small, VGG) for the classification of SAT-4/SAT-6 dataset (from US National Agriculture Imagery Program, NAIP) using Convolutional Neural Networks (CNN). Pirotti et al. [4] benchmarked different machine learning methods (including multilayered perceptron) for classification of Sentinel-2 images. In [5], we proposed a new 3D CNN architecture for hyperspectral data pixelwise classification (semantic segmentation). In [6, 7], we adapted the SegNet architecture to achieve semantic segmentation of multimodal airborne imagery.

In this paper, we rely on the recent DenseNet [8] and SegNet [9] architectures to perform land cover semantic segmentation of large multispectral Sentinel-2 images. These architectures are experimentally assessed through two different use cases, namely fine and coarse resolution estimation. We also introduce a new 3D DenseNet network in order to jointly process both spatial and spectral dimensions. In addition, we suggest the use of a “noisy ground truth” (i.e. outdated and low spatial resolution labels) for both training and testing. The idea is to explore the feasibility of outdated low quality knowledge integration for modern image sensors and analysis methods. Experiments are conducted in a wide region between France, Switzerland and Italy relying on the reference areas of GlobCover (ESA 2009 Global Land Cover Map) annotation and 2016 ESA Sentinel-2 images.

## 2. SPECTRAL CHANNEL FUSION AND DEEP NEURAL NETWORKS


When dealing with multispectral images, strategies for processing and fusing spectral channels are numerous. Deep neural networks enable a large variety of choices from the

component operator to architectural levels. At the low component level, the classical approach is to use 2D convolution layers from the first stages of the network. In such a configuration, each neuron applies a specific filter to each channel and then fuses (sums) the resulting maps. Such an approach enables the early combination of each channel of multispectral image sensors whatever their wavelength. One can also rely on 3D convolution neurons that consider input images as 3D volumes. These volumes are then filtered by 3D filters that locally combine the spatial and spectral information. With such a strategy, cascading 3D filters enables all the available wavelength bands to be combined in a growing bandwidth manner. From a coarse network architecture point of view, the advances in neural network modeling led to an increase in architecture depth and a variety of neural layers combination enabling data representation levels to be also fused together.

In classical approaches such as AlexNet [10], a neuron layer is only fed by the previous layer output such that low-level representations are not directly included at the decision level that occurs in the last layer. However, recent ideas enable a given neuron layer to fuse the information coming from many more previous neural layers. The Residual Network (**ResNet**) components [11] consist in the fusion (pixelwise sum) between the resulting representation of a block of 2D convolution filters and its input. This enables to mix input data with a new slightly increased representation level.

One step further, the **SegNet** architecture [9] for semantic segmentation adopts an auto-encoder structure, with an encoder that models the data at multiple scales and a decoder that up-samples the internal representation and projects it into the mapping space. To do so, the location of the locally maximal activation from the lower layers are fed forward directly to their up-sampling counterpart in the last layers, which allows the decoder to relocate abstract features into the most salient points originally detected. This information is actually crucial for accurate semantic map boundary reconstruction.

Even further, the recent **DenseNet** architecture [8] proposes intensive layer fusion. Given a *DenseBlock* that consists in a set of convolution layers working at the same scale, each neural layer processes the concatenation of all its previous layers thus enabling the fusion of very numerous representation levels. Similarly to SegNet, its semantic segmentation extension [12] adds a decoding path to generate the semantic map. However, on the decoding branch, the fusion not only consists in intra dense block layers fusion but it also relies, at the input of each decoding block, on the concatenation of the preceding high level feature maps and the ones coming from the encoding block at the same scale. The strength of such an approach is the intensive feature reuse all along the network thus enabling the number of parameters to be significantly reduced while increasing performance.



**Fig. 1: Top:** reference DenseNet architectures. **Bottom:** our proposed 3D DenseNet

### 3. PROPOSED APPROACHES

In this paper we propose to explore the capabilities of some recent neural networks discussed in Sec. 2 for land cover mapping. Conversely to public datasets provided with recent contests such as the semantic labeling benchmark over ISPRS Vaihingen and Potsdam cities that relies on high-quality ground truth, we propose to train deep networks on a wide area, relying on coarse and noisy reference data. This realistic scenario allows to exploit previous knowledge such as outdated and coarse resolution land cover maps. These maps bring both confident annotations on qualitative and temporally stable regions, while presenting erroneous values on area boundaries but also on areas where changes occurred before image acquisition. In this context, training neural networks in a supervised way is a delicate step and we propose two different use cases considering the following architectures, whose results are reported in Tab.1.

First, we consider fine-grained pixel-level land cover estimation using very light architectures with few parameters using DenseNet [8]. 3 basic configurations are proposed, with either 2 or 3 dense blocks at the encoder and decoder levels, different number of layers per block and either 12 or 16 neurons per layer. As shown in Fig. 1, an extension of this architecture is proposed by introducing 3D convolution neurons in the first dense block. As already proposed in [5], we aim to early combine spectral channels but only for adjacent bands. To do so, the  $48 \times 2D$  convolution neurons input layer of the reference model is removed thus enabling all the input spectral bands to feed a new first dense block exclusively composed of 3D convolutions with kernel size  $3 \times 3 \times 3$  using 4 layers of 16 neurons each. Then, if 9 spectral bands regularly sampled along the wavelength dimension are chosen, the neurons of the 4th layer of the 3D dense block have processed data that depend on the full spectral bandwidth. An intermediate 3D convolution is then added next to ensure the 3D to 2D data transition by squeezing the spectral dimension. Com-

pared to 2D convolution based models, the overall number of parameters decreases by at least 10%.

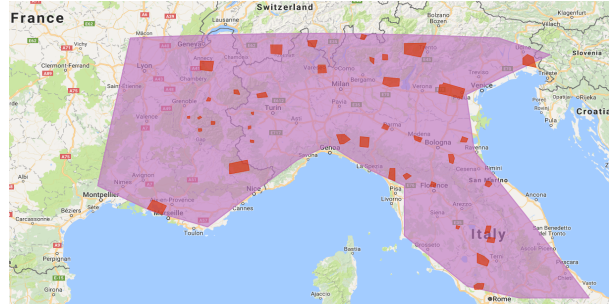
Second, we consider land cover estimation at the coarse ground truth scale using a very large architecture derived from SegNet [9]. We modify the original SegNet by allowing the input to have more than 3 bands by enlarging the first convolution layer. The encoder is based on VGG-16 [13], which alternates blocks of 2 or 3 convolutions with  $3 \times 3$  kernels with max-pooling layer that downsamples the data by a factor of two. The decoder is symmetric to the encoder, with sparse  $2 \times$  up-sampling instead of pooling and convolutions to densify the activations. The up-sampling is performed using the locations of the maxima found during pooling in the encoder, that are skipped through the network. After each block in the decoder, a  $1 \times 1$  convolution layer projects the feature maps into the label space. We add decoding blocks until the estimation has a higher resolution than the ground truth. The multiscale estimations are then interpolated at full scale and averaged for loss computation. The gradient is then back-propagated in a deeply supervised fashion, to enforce a better multiscale spatial regularity. The full architecture has more than 27M parameters and is the largest model presented in this work.

#### 4. EXPERIMENTS

The proposed neural networks are trained and evaluated on a dataset extracted from a region between France, Switzerland and Italy as shown in Fig. 2. Images have been acquired by the Sentinel-2 sensor within the May-October 2016 period, while land cover ground truth is the 2009 GlobCover map *ESA/GLOBCOVER\_L4.200901\_200912.V2.3*. The Google Earth Engine [14] was used to extract Sentinel Images but only on the reference areas of GlobCover. Some subregions of this dataset have been put apart to serve as the test dataset. Two datasets are proposed and detailed in Tab.2: the first one covers all the May-October period and does not include areas with clouds (opaque and cirrus). The second one only covers the summer period but includes clouds so that they can be considered as an additional class to detect. Both datasets are then multitemporal with frequent observations of the same areas at different timestamps.

The Sentinel-2 data is interpolated to 20m/px resolution, while the GlobCover ground truth is at 300m/px resolution. The proposed networks are trained considering two strategies: the light ones are optimized at the image resolution level whatever the ground truth resolution is (*i.e.* the models are trained against an interpolated ground truth), thus trying to estimate at the scale of the input even if the reference is too coarse. The largest network (SegNet) however estimates multiple maps at several scales that are averaged and interpolated to the ground truth resolution. Training at a lower resolution alleviate errors along class borders in the GlobCover data.

As shown on Tab. 1, estimation at a higher resolution than



**Fig. 2:** Experimental datasets, training images were extracted from the overall pink region while test regions are the small orange ones. Each image covers GlobCover reference areas.

the reference provides an overall accuracy ranging between 25.1% and 30.0% on dataset D1 and 39.5% and 57.9% on D2. However, the overall accuracy metric introduces a discrepancy that penalizes the reported values. The metric is actually confident on stable and large areas but is less confident on heterogeneous areas. In such a noisy training context, resembling classes such as “Rainfed croplands” and “Mosaic Cropland/Vegetation” are often confused. However, these network are expressive enough to well detect classes such as “Artificial surfaces”, “Bare areas”, “Water bodies” and the added class “Clouds”. When working with only the regularly sampled spectral bands (9 bands tests), 3D DenseNets can be applied. 3D approaches are performing worse but performance increases with model depth. Nevertheless, many more training images are required to train for the fine-grained approach. However, coarse land cover estimation provides good results on both datasets. The large number of parameters of SegNet and the multiscale approach provides at least 66.9% accuracy on the long period dataset that shows strong aspect changes of the same areas (D1) and 83.9% on the more stable one (D2).

#### 5. CONCLUSION

This paper reviews a variety of strategies provided by Deep Learning approaches to fuse the channels of multispectral sensors with the aim of land cover mapping. Optimizing such neural networks from a noisy land cover reference is challenging when dealing with reduced size datasets. Estimating at coarse image scale up to the one of the reference remains the most appropriate solution. Estimating at a higher resolution is difficult for the training step but also at the validation step to enable confident comparison. A first step forward would consist in enhancing the proposed models by enabling multiscale estimation for each of them. Going further, estimating at a finer resolution is challenging but required for land cover monitoring, and refined approaches should be investigated by taking advantage of the recently available mass of training data.

**Table 1:** Description of the considered architectures and Overall Accuracy (OA) after reference class boundary erosion of 200m. For DenseNet, e[x,y], d[y,x] and b[z] list the number of blocks and their respective number of layers for respectively the encoding and symmetric decoding branches and the bottleneck block. The number of neurons per layer is parametrized by  $g$ .

Use Case	Models	#params	#scales	OA/9 bands [B1-B8a]		OA/13 bands	
				D1	D2	D1	D2
Fine-grained estimation	DenseNet						
	e[2,3],b[4],d[3,2] g12	0.23M	3	30.1%	57.9%	35.9%	51.8%
	e[4,5],b[7],d[5,4] g16	1.00M	3	29.5%	55.3%	27.2%	55.4%
	e[4,4,4],b[4],d[4,4,4] g16	1.08M	4	25.7%	52.3%	28.5%	51.4%
Fine-grained estimation	3D DenseNet						
	e[4, 5] <sub>3D</sub> ,b[7],d[5,4] g16	0.88M	3	25.1%	39.5%	-	-
	e[4, 4, 4] <sub>3D</sub> ,b[4],d[4,4,4] g16	0.92M	4	26.5%	41.2%	-	-
Coarse estimation	SegNet	27.00M	5	-	-	66.9%	83.9%

**Table 2:** Considered datasets: images including clouds detected by the Sentinel-2 A60 additional band are excluded from the long period dataset (D1). The short period dataset (D2) includes clouds and enables training for their detection. Both datasets collect more than 150M pixels each.

Dataset (period)	#images		
	train	test	#class
D1, Long period no clouds (May-Oct. 2016)	140	54	23
D2, Short period with clouds (June-August 2016)	158	39	24

## 6. REFERENCES

- [1] R. Bokusheva, F. Kogan, I. Vitkovskaya, S. Conradt, and M. Batyrbayeva, “Satellite-based vegetation health indices as a criteria for insuring against drought-related yield losses,” *Agricultural and Forest Meteorology*, vol. 220, no. 15, pp. 200–206, 2016.
- [2] ESA, *Satellite Earth Observations in Support of Disaster Risk Reduction*, Special 2015 Edition for the World Conference on Disaster Risk Reduction, 2015.
- [3] M. Papadomanolaki, M. Vakalopoulou, S. Zagoruyko, and K. Karantzalos, “Benchmarking deep learning frameworks for the classification of very high resolution satellite multispectral data,” *ISPRS Annals of Photogrammetry, Remote Sensing and Spatial Information Sciences*, vol. III-7, pp. 83–88, 2016.
- [4] F. Pirotti, F. Sunar, and M. Piragnolo, “Benchmark of Machine Learning Methods for Classification of a SENTINEL-2 Image,” *ISPRS - International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences*, pp. 335–340, June 2016.
- [5] A. Ben Hamida, A. Benoit, P. Lambert, and C. Ben Amar, “Deep learning approach for remote sensing image analysis,” in *Conference on Big Data from Space*. BIDS, 2016.
- [6] N. Audebert, B. Le Saux, and S. Lefèvre, “Semantic segmentation of earth observation data using multi-modal and multi-scale deep networks,” in *Asian Conference on Computer Vision*, 2016.
- [7] N. Audebert, B. Le Saux, and S. Lefèvre, “Fusion of heterogeneous data in convolutional networks for urban semantic labeling,” in *Joint Urban Remote Sensing Event*, 2017.
- [8] M. Volpi and D. Tuia, “Dense semantic labeling of sub-decimeter resolution images with convolutional neural networks,” *IEEE Transactions on Geoscience and Remote Sensing*, 2016.
- [9] V. Badrinarayanan, A. Kendall, and R. Cipolla, “SegNet: A Deep Convolutional Encoder-Decoder Architecture for Scene Segmentation,” *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. PP, no. 99, pp. 1–1, 2017.
- [10] A. Krizhevsky, I. Sutskever, and G.E. Hinton, “Imagenet classification with deep convolutional neural networks,” in *Advances in neural information processing systems*, 2012, pp. 1097–1105.
- [11] K. He, X. Zhang, S. Ren, and J. Sun, “Deep residual learning for image recognition,” *CoRR*, vol. abs/1512.03385, 2015.
- [12] S. Jégou, M. Drozdal, D. Vázquez, A. Romero, and Y. Bengio, “The one hundred layers tiramisu: Fully convolutional densenets for semantic segmentation,” *CoRR*, vol. abs/1611.09326, 2016.
- [13] K. Chatfield, K. Simonyan, A. Vedaldi, and A. Zisserman, “Return of the devil in the details: Delving deep into convolutional networks,” August 2014.
- [14] Google Earth Engine Team, “Google earth engine: A planetary-scale geo-spatial analysis platform,” <https://earthengine.google.com>, 12.