

Uncertainty Gated Network for Land Cover Segmentation

Guillem Pascual
University of Barcelona
guillem.pascual@ub.edu

Santi Seguí
University of Barcelona
santi.segui@ub.edu

Jordi Vitrià
University of Barcelona
jordi.vitria@ub.edu

Abstract

The production of thematic maps depicting land cover is one of the most common applications of remote sensing. To this end, several semantic segmentation approaches, based on deep learning, have been proposed in the literature, but land cover segmentation is still considered an open problem due to some specific problems related to remote sensing imaging. In this paper we propose a novel approach to deal with the problem of modelling multiscale contexts surrounding pixels of different land cover categories. The approach leverages the computation of a heteroscedastic measure of uncertainty when classifying individual pixels in an image. This classification uncertainty measure is used to define a set of memory gates between layers that allow a principled method to select the optimal decision for each pixel.

1. Introduction

Land cover segmentation deals with the problem of multi-class semantic segmentation of remote sensing images. This problem, which consists of assigning a unique label (or class) to every pixel of an image, is particularly difficult due to (i) the high resolution of the images and diversity of size of the objects, (ii) the diversity of classes and, usually, the similarities among them, (iii) the noisy labeling and implicit rules such as not considering small/isolated areas and (iv) data domain: the model is usually trained with a set of images that highly differ from the target area where it is expected to generalize and perform predictions.

In the recent literature, most of the methods solving these problems are based on deep learning. In [6] Long *et al.* popularized the use of fully convolutional networks for segmentation. This method without any dense layer, allowed to create segmentation maps for images of any size. Based on this idea, and also trying to solve the exact alignment problem associated with the pooling layers, several methods have been presented [2, 8, 12, 4]. U-Net [8] is a popular architecture defined as an encoder-decoder scheme, where in the encoder stage the spatial dimension is gradually reduced with pooling layers and then decoder stage gradually




Figure 1: Prediction examples from the model. Each row is a sample, the left column is the input image and the right column is the predicted segmentation map.

recovers the object details and spatial dimension to finally obtain the output segmentation map. In RefineNet [4], proposed by Lin *et al.*, the ResNet architecture is used as an encoder step while in the decoder step as a set of RefineNet blocks which fuse high resolution features from the encoder and low resolution features from previous RefineNet block.

In the domain of satellite images, several methods trying to solve this problem in high-resolution images have been presented [9, 10, 7, 5]. The most relevant publication for our work is Gated Convolutional Network (GCN) [10], where the segmentation is computed from the outputs of each block of a pre-trained ResNet, using entropy as a gate to fine-tune the prediction at each level.

In this paper, we propose a novel method that tackles the problem of land cover segmentation using the data and protocol proposed by the DeepGlobe Land Cover Classification Challenge at CVPRW [1]. Figure 1 shows two samples from the dataset and the predictions of our model. The proposed method is built over a GCN using a ResNet architecture and exploits the uncertainty of the predictions in each layer. The uncertainty measure, built on the basis of

the publication by Alex Kendall and Yarin Gal [3], is used to define a set of memory gates between layers that allow for a principled method to select the optimal decision for each pixel.

The remainder of this paper is organized as follows. In the next section we present the proposed method. In Section III, we present the training setup. In Section IV, we present the experimental results. Finally, Section IV concludes the paper with remarks on the proposed approach.

2. Method

Our model builds upon the GCN architecture proposed by Wang *et al.* in [10]. In that paper a new architecture was proposed to combine the feature maps learned at different blocks of a ResNet model by using memory gates instead of more classical operations such as summation or concatenation. The gating mechanism was based on the relationship between the information entropy of the feature maps and the label-error map, allowing for a better feature map integration.

To further develop the concept of gated convolutions, we consider the use of a more principled concept: assigning a credibility measure to each feature map. According to the Bayesian viewpoint proposed by Alex Kendall and Yarin Gal in [3], it is possible to characterize the concept of uncertainty into two categories. On the one hand, if the noise applies to the model parameters, we will refer to *epistemic uncertainty*. On the other hand, if the noise occurs directly in the observation, we will refer to it as *aleatoric uncertainty*. Additionally, aleatoric uncertainty can further be categorized into two more categories: *homoscedastic uncertainty*, when the noise is constant for all the outputs (thus acting as a “measurement error”), or *heteroscedastic uncertainty* when the noise of the output also depends explicitly on the specific input.

We propose to use a measure of heteroscedastic uncertainty when classifying specific pixels as a gating mechanism. In this case, we have to measure the heteroscedastic uncertainty in a classification task, where the noise model is placed in the *logit* space.

Let σ_i and l_i be two predicted vectors of unaries of dimension C , the number of classes, for every input pixel x_i . The latter, l_i , are the logits used to output a probability distribution by using a softmax, while σ_i aims to bound its uncertainty. By taking T random samples of l_i perturbed by σ_i , we can derive a stochastic loss $\mathcal{L}_x$ that allows the computation of an uncertainty value γ_i for each x_i input as follows:

$$\mathcal{L}_x = \sum_i \gamma_i$$

$$\gamma_i = -\log \frac{1}{T} \sum_t \exp(\hat{l}_{i,t,c} - \log \sum_{c'} \exp \hat{l}_{i,t,c'})$$

$$\hat{l}_{i,t} \sim \mathcal{N}(l_i, \sigma_i), 1 \leq t \leq T$$

where $\hat{l}_{i,t,c'}$ is the t sampled logit vector from class c' , and $\hat{l}_{i,t,c}$ is the logit vector of the winner class for each pixel and sample.

Our architecture is illustrated in Figure 2. As it can be seen, an uncertainty measure $\gamma^{(j)}$ is computed after each of the ResNet blocks g_j , $0 \leq j \leq 4$. The blocks g_4 through g_1 correspond to each of the original residual blocks, while g_0 is composed by the first max-pooling and convolution.

The refinement process through uncertainty gates starts by setting $\bar{b}_4 = g_4$ and b_j the upsampled version of $\bar{b}_j$ to match g_{j-1} dimensions. Then for each $j = 4, \dots, 1$ the process of obtaining an uncertainty and segmentation is defined as follows:

$$l^{(j)} = b^{(j)} \otimes_c \mathbf{w}_{1 \times 1}^{(j,1)}$$

$$\sigma^{(j)} = b^{(j)} \otimes_c \mathbf{w}_{1 \times 1}^{(j,2)}$$

$$\hat{l}_{i,t}^{(j)} \sim \mathcal{N}(l_i^{(j)}, \sigma_i^{(j)})$$

$$\gamma^{(j)} = \log \frac{1}{T} \sum_t \exp(\hat{l}_{i,t,c}^{(j)} - \log \sum_{c'} \exp \hat{l}_{i,t,c'}^{(j)})$$

$$\bar{b}_{j-1} = \gamma^{(j)} * g_{j-1} + b_j$$

Where $\otimes_c$ is the convolution operator with a 1×1 kernel and dimension C , and $*$ denotes the element-wise multiplication, but defined in such a way that gradient can only flow through the $\bar{g}_j$ operand during the backpropagation step. If gradient is allowed to flow through $\gamma^{(j)}$ in the backward pass, we can no longer talk about heteroscedastic uncertainty, as external factors aside from pure classification would condition them. Finally, $\gamma^{(0)}$ and $L^{(0)}$ is computed in the same manner.

To compute the final segmentation, instead of taking new logits from the last block as GCN does, it is proposed a method that takes advantage of all logits $l^{(j)}$ and uncertainties $\gamma^{(j)}$ calculated at each block. The final probabilities of each pixel and class are obtained with a γ -weighted sum of the probabilities at each intermediate step, as:

$$\frac{1}{C} \sum_{i=0}^4 \text{softmax}(l^{(i)}) * (1 - \gamma^{(i)})$$

3. Training

To train the model, we first reduce the original resolution down to 1024×1024 , which simplifies the problem space while keeping enough details. The network is further trained by taking 8 random crops, each of 250×250 , out of each image. Each crop is then randomly rotated and flipped, and is further processed by adding gaussian noise and adjusting hue, contrast and brightness.




Figure 2: Uncertainty gated convolutional neural network. Black arrows represent weighted connections between different layers. Green arrows represent forward-only weighted connections, where gradient flows in the backpropagation process are not allowed.

The model is trained by minimizing, at each level j , both $L^{(j)}$ and a classification loss given by a softmax cross-entropy between the labels and sampled unaries from the logits. Overall, loss is minimized with WNAdam optimizer [11], using an standard piecewise learning rate decay for a total of 100 epochs.

4. Results

The data for the DeepGlobe Land Cover Classification Challenge consists of 1.146 satellite RGB images of size 2448×2448 pixels, split into training/validation/test, each with 803/171/172 images. Each satellite image is paired with a class labeled image using the following 7 categories: 1) Urban land; 2) Agriculture land; 3) Rangeland; 4) Water (Rivers, oceans, lakes, wetland, ponds); 5) Barren land (Mountain, land, rock, dessert, beach, no vegetation) and 7) Unknown (clouds and other artifacts).

The pixel-wise mean Intersection over Union (mIoU) score, calculated by averaging the IoU over all classes, is used as evaluation metric. The IoU is defined as: True Positive / (True Positive + False Positive + False Negative). The *unknown* class is not an active class used in evaluation.

The final model uses, as discussed in the previous section, the ResNet with 18 layers and is trained for 100 epochs. Figure 3 shows the prediction process. Segmen-

tations at each level are generated for visualization and interpretation and further combined to obtain the final result. Deeper levels are more general and can not accurately predict each pixel, which can be both attributed to the down-sampling process and the abstraction done through all the convolutions. That is why upper levels refine the result and are richer in details. In particular, it can be seen that pixels where the output does not match the ground truth, a high uncertainty is obtained. Averaging across all levels improves the result by reducing artifacts and producing smoother segmentation maps. The model runs inference in real time, taking only 250ms to produce a segmentation at full 2448×2448 resolution on an NVIDIA Titan X. This architecture achieves a mIoU score of 0.485 in the final test set of the challenge.

5. Conclusions

In this paper, an uncertainty gated convolutional neural network has been proposed for land-cover semantic segmentation. The proposed method leverages the computation of a heteroscedastic measure of uncertainty when classifying individual pixels in an image. This classification uncertainty measure is used to define a set of memory gates between layers that allow for a principled method to select the optimal decision for each pixel. The result reported on the DeepGlobe Land Cover Classification Challenge is 0.485 mIoU on the final test set. Future improvement on the domain adaptation problem will be considered, since we have observed some inconsistencies due to this specific issue.

Acknowledgements

This work was partially founded by MINECO Grant TIN2015-66951-C2 and by an FPU grant (Formacin de Profesorado Universitario) from the Spanish Ministry of Education, Culture and Sport (MECD) to Guillem Pascual (FPU16/06843). We gratefully acknowledge the support of NVIDIA Corporation with the donation of the Titan X Pascal GPU used for this research.

References

- [1] I. Demir, K. Koperski, D. Lindenbaum, G. Pang, J. Huang, S. Basu, F. Hughes, D. Tuia, and R. Raskar. Deepglobe 2018: A challenge to parse the earth through satellite images. *arXiv preprint arXiv:1805.06561*, 2018. 1
- [2] S. Jgou, M. Drozdal, D. Vazquez, A. Romero, and Y. Bengio. The one hundred layers tiramisu: Fully convolutional densenets for semantic segmentation. In *2017 IEEE Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)*, pages 1175–1183, July 2017. 1
- [3] A. Kendall and Y. Gal. What uncertainties do we need in bayesian deep learning for computer vision? In *Advances in*




Figure 3: Prediction process of an image. All b_j are intermediate segmentation outputs, resulting from a $\text{softmax}(l^j)$, and upsampled to target resolution. Low uncertainty is represented by blue, yellow indicates high uncertainty.

Neural Information Processing Systems, pages 5580–5590, 2017. [2](#)

- [4] G. Lin, A. Milan, C. Shen, and I. Reid. Refinenet: Multi-path refinement networks for high-resolution semantic segmentation. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017. [1](#)
- [5] Y. Liu, D. Minh Nguyen, N. Deligiannis, W. Ding, and A. Munteanu. Hourglass-shapenetwork based semantic segmentation for high resolution aerial imagery. *Remote Sensing*, 9(6), 2017. [1](#)
- [6] J. Long, E. Shelhamer, and T. Darrell. Fully convolutional networks for semantic segmentation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3431–3440, 2015. [1](#)
- [7] K. Nogueira, M. Dalla Mura, J. Chanussot, W. Robson Schwartz, and J. A. dos Santos. Dynamic Multi-Scale Semantic Segmentation based on Dilated Convolutional Networks. *ArXiv e-prints*, Apr. 2018. [1](#)
- [8] O. Ronneberger, P. Fischer, and T. Brox. U-net: Convolutional networks for biomedical image segmentation. In N. Navab, J. Hornegger, W. M. Wells, and A. F. Frangi, editors, *Medical Image Computing and Computer-Assisted Intervention – MICCAI 2015*, pages 234–241, Cham, 2015. Springer International Publishing. [1](#)
- [9] J. Sherrah. Fully convolutional networks for dense semantic labelling of high-resolution aerial imagery. *CoRR*, abs/1606.02585, 2016. [1](#)
- [10] H. Wang, Y. Wang, Q. Zhang, S. Xiang, and C. Pan. Gated convolutional neural network for semantic segmentation in high-resolution images. *Remote Sensing*, 9(5), 2017. [1](#), [2](#)
- [11] X. Wu, R. Ward, and L. Bottou. Wngrad: Learn the learning rate in gradient descent. *arXiv preprint arXiv:1803.02865*, 2018. [3](#)
- [12] H. Zhao, J. Shi, X. Qi, X. Wang, and J. Jia. Pyramid scene parsing network. In *Proceedings of IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017. [1](#)