



HAL
open science

Weakly Supervised Semantic Segmentation of Satellite Images

Adrien Nivaggioli, Hicham Randrianarivo

► **To cite this version:**

Adrien Nivaggioli, Hicham Randrianarivo. Weakly Supervised Semantic Segmentation of Satellite Images. Joint Urban Remote Sensing Event (JURSE), May 2019, Vannes, France. hal-02091519

HAL Id: hal-02091519

<https://hal.science/hal-02091519>

Submitted on 5 Apr 2019

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Weakly Supervised Semantic Segmentation of Satellite Images

Adrien Nivaggioli
Qwant Research
Paris, France
a.nivaggioli@qwant.com

Hicham Randrianarivo
Qwant Research
Paris, France
h.randrianarivo@qwant.com

Abstract—When one wants to train a neural network to perform semantic segmentation, creating pixel-level annotations for each of the images in the database is a tedious task. If he works with aerial or satellite images, which are usually very large, it is even worse. With that in mind, we investigate how to use image-level annotations in order to perform semantic segmentation. Image-level annotations are much less expensive to acquire than pixel-level annotations, but we lose a lot of information for the training of the model. From the annotations of the images, the model must find by itself how to classify the different regions of the image. In this work, we use the method proposed by Anh and Kwak [1] to produce pixel-level annotation from image level annotation.

We compare the overall quality of our generated dataset with the original dataset.

In addition, we propose an adaptation of the AffinityNet that allows us to directly perform a semantic segmentation.

Our results show that the generated labels lead to the same performances for the training of several segmentation networks. Also, the quality of semantic segmentation performed directly by the AffinityNet and the Random Walk is close to the one of the best fully-supervised approaches.

Index Terms—Computer vision, Weak learning, Semantic segmentation, Land cover classification

I. INTRODUCTION

Semantic segmentation of satellite and aerial images could be incredibly helpful for fields like urban planning, disaster recovery, autonomous agriculture, environmental monitoring and many others. We now have access to large databases filled with more images than any manual method could handle (such as the USGS Earth Explorer¹, ESAs Sentinel Mission² or NASA’s Earthdata Search³). The need for an automatized study of those images is obvious, and there is an urgent demand for tools and methods that allow automatic interpretation of this huge amount of data. In the last few years, deep learning has become the essential tool for solving this kind of problem [8], [10]. For the task of semantic segmentation, several successful methods appeared recently like SegNet [3], FCN [18], U-NET [17] or PSPNet [22] with great results in satellite and aerial imagery [2], [13], [19]

One of the biggest difficulties faced by those methods is the lack of pixel-level labels for those images. Indeed, the

process of manually annotating each image is tedious. This is also an issue when trying to perform a semantic segmentation of a regular image, like those of the Pascal VOC [7] or the Common Object in COntext (COCO) datasets [12]. To face it, methods relying on weaker types of labels started to appear, called Weakly Supervised Learning (WSL). Those methods can use simple bounding-box annotations [9], [15], scribbles [11], [20], points [5], or a simple image-level class label [14], [16].

The method we use is divided into 4 steps (c.f. fig. 1), and based on the work of Ahn and Kwak [1]. First, a classification network is trained using image level annotations. Then, a second network is trained to learn the relationships between a pixel and its neighbourhood, called an affinity network. After this, a random walk is performed combining Class Activation Map (CAM) and affinity labels to produce the segmentation labels. Finally, we can use the segmentation labels produced by the method to train a segmentation model.

Contributions: We use a new backbone better adapted to satellite imagery. We conduct several experiments on the loss function proposed by [1] to adapt its hyperparameters for satellite and aerial imagery. We train several state-of-the-art segmentation models (SegNet [3], PSPNet [22] and UNet [17]) to validate the quality of the generated labels for the training of segmentation models. We modify the method in order to perform semantic segmentation using only the classification and the affinity networks. (c.f. fig. 1) Finally, we compare our semantic segmentation results with fully supervised approaches on the validation set of the DEEPGLOBE dataset [6] and study the trade-off between information and quality of segmentation..

II. METHOD

In this section, we present the approach we follow for weakly based semantic segmentation. The method can be split into 4 different parts as shown in fig. 1 : Classification Network, Affinity Network, Random Walk and Segmentation Network.

A. Classification Network

The first step of the method is the training of a classification network that will be used to identify the different categories present in the images. The main idea is to use the capacity

¹<https://earthexplorer.usgs.gov/>

²<https://sentinel.esa.int/web/sentinel/home>

³<https://search.earthdata.nasa.gov/>

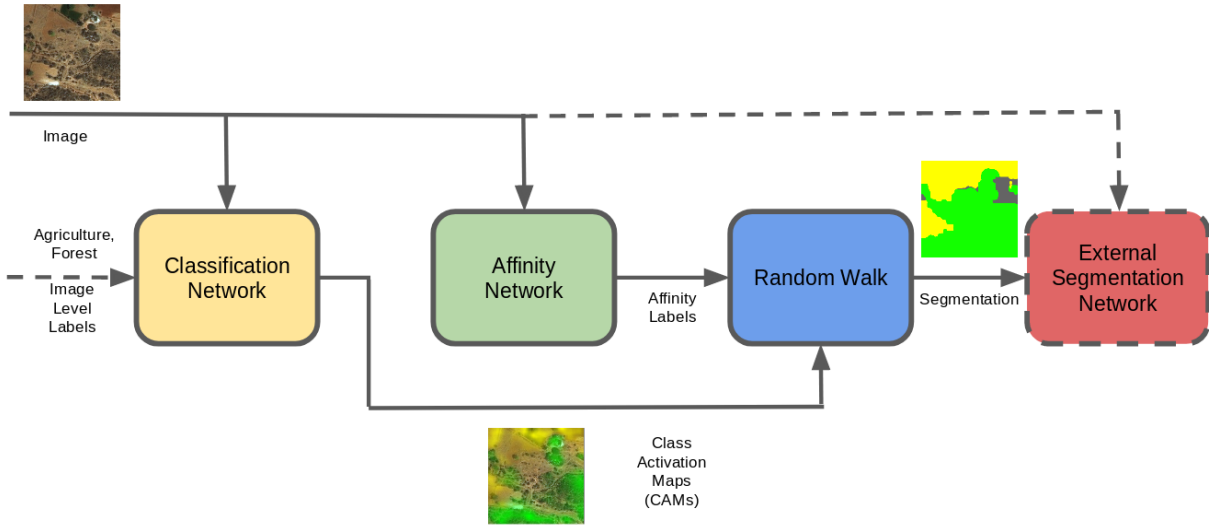


Fig. 1: Inference pipeline of the original Affinity-Net, that requires both an image and a label as inputs. We propose to remove the dotted parts to transform the Affinity-Net into an independent segmentation network.

of the CAM that can be extracted from the Convolutional Neural Network (CNN) to localize in the image the areas that influence the most the classification result [4], [16], [23]. For each image, the classification network will produce a set of CAMs $\mathcal{M} = \{M_c \mid c \in \mathcal{C}\}$, each CAM M_c corresponding to the activation map of the class c . Additionally, another CAM M_{bg} is defined to localize background in the image (1).

$$M_{bg}(x, y) = \{1 - \max_{c \in \mathcal{C}} M_c(x, y)\}^\alpha \quad (1)$$

Usually, there is no background in aerial images, but this CAM prevents the network to predict unsure result. Even if this means that a smaller portion of the original dataset will be usable to train the segmentation network, we favour precision over quantity. With that in mind, we conducted all of our experiments both with and without the background, in order to compare the results and assess which approach fits our case best.

B. Affinity Network

The second part of the method consists in training an Affinity Network which models relationships between pairs of pixels (i, j) in the image. The AffinityNet is designed to extract a convolutional feature map f^{aff} where each element can be seen as an affinity feature. The affinity between two pixels (i, j) denoted by W_{ij} is defined as the similarity between the affinity features:

$$W_{ij} = \exp\{-\|f^{\text{aff}}(x_i, y_i) - f^{\text{aff}}(x_j, y_j)\|_1\} \quad (2)$$

Nevertheless, the affinity labels needed to train the AffinityNet are not directly available. A way to generate these affinity labels is to use the CAM as partial sources of supervision for their generation. The values of the CAM are used as a confidence score to determine whether a pixel belongs to a

category c or not. A high α is used to assess confident regions of classes $c \in \mathcal{C}$, and a lower one is used to assess areas the network is the most unsure about. The pixels that are left are considered as neutral. Using the CAM with high confidence, a binary label is given to the pair of pixels. If their classes are the same, their affinity label W_{ij}^* is 1, 0 otherwise. The pairs containing neutral pixel are considered neutral and ignored during the training (c.f. fig. 2).

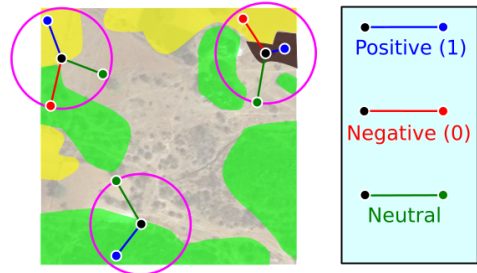


Fig. 2: Generating Affinity labels. The areas correspond to confident prediction of *Agriculture* in yellow and *Forest* in green. The black area is the *background* and the white is *neutral*.

From there, we compute for each coordinate of the image its affinities with the other coordinates within a circle of fixed radius γ . The set of pairs of coordinates that have an affinity is defined by :

$$\mathcal{P} = \{(i, j) \mid d((x_i, y_i), (x_j, y_j)) < \gamma, \forall i \neq j\} \quad (3)$$

The set \mathcal{P} is then divided into two subset \mathcal{P}^+ and \mathcal{P}^- where $W_{ij}^* = 1$ and $W_{ij}^* = 0$ respectively. The subset \mathcal{P}^+ is again divided into two subsets \mathcal{P}_{fg}^+ for foreground and \mathcal{P}_{bg}^+ for the background regions.

The model is trained using the following loss function :

$$\mathcal{L} = \frac{\mathcal{L}_{fg}^+}{a} + \frac{\mathcal{L}_{bg}^+}{b} + \frac{\mathcal{L}^-}{c} \quad (4)$$

with

$$\frac{1}{a} + \frac{1}{b} + \frac{1}{c} = 1 \quad (5)$$

$$\mathcal{L}_{fg}^+ = -\frac{1}{|\mathcal{P}_{fg}^+|} \sum_{(i,j) \in \mathcal{P}_{fg}^+} \log W_{i,j} \quad (6)$$

$$\mathcal{L}_{bg}^+ = -\frac{1}{|\mathcal{P}_{bg}^+|} \sum_{(i,j) \in \mathcal{P}_{bg}^+} \log W_{i,j} \quad (7)$$

$$\mathcal{L}^- = -\frac{1}{|\mathcal{P}^-|} \sum_{(i,j) \in \mathcal{P}^-} \log (1 - W_{i,j}) \quad (8)$$

This loss is quite similar to the one used by the authors of [1], but they suggest that a should be equal to b . We want to be able to penalize the background further, so we slackened the constraints. [1] used eq. 4 with $a = b = 4, c = 2$. In order to penalize the background further, we used $a = 6, b = 2, c = 3$.

Those labels can then be used to perform a random walk on the original CAMs in order to get proper segmentation labels.

From there, the original method uses those labels to train a regular segmentation network. We compared the results of 3 different networks, U-Net [17], PSPNet [22] and SegNet [3] using our generated labels.

We also tried a new approach that allows us to directly performs semantic segmentation with the trained classification and affinity networks. Indeed, [1] proposes to set to 0 the CAMs of the class we know are not present in the image. While this increases the quality of the segmentation labels, it prevents us from using the AffinityNet on images we do not have any labels on. To assess which CAM need to be kept, we use the confidence scores output by the classification network.

As a backbone network, we used an extension of the ResNet38 [21], composed of 74 convolutional layers (as opposed to 38 in the ResNet38). We found that a deeper network greatly improved our performances, supposedly because aerial images have a lower variance than regular ones, thus making it more difficult for the network to differentiate the different classes.

III. EXPERIMENTS AND RESULTS

A. Dataset

We used the DeepGlobe dataset [6] composed of 803 aerial images with pixel-level labels. We kept 562 for training and the remaining 241 for validation. Each image has a size of 2448x2448 pixels, and contains at least one of those categories : Barren, Water, Urban, Forest, Agriculture, Rangeland and Unknown. We split each image into 64 patches of size 306x306 each. Even if the dataset comes with pixel-wise annotations of the images, we reduced those labels to image-level (which categories are present in the labels). This means that no information about localization and/or distribution of

the categories are present in the final dataset. The DeepGlobe dataset also offers 171 aerial images without any labels that we used for testing. Indeed, we can upload our results to their CodaLab competition⁴ and compare our score to others, all of them using fully-supervised techniques.

B. Quality of Segmentation Labels

We used the training dataset to teach both the classification Network and the Affinity Network. Then, we generated the Segmentation Label from the validation dataset. Because we have the pixel-level labels of those images (even if they were not used during the training), we can compare the quality of our segmentation labels to the ground truth provided by DeepGlobe as shown in fig. 3. [1] showed that their method was fairly insensitive to hyperparameters, but our dataset of aerial images is different from regular images, mostly because they have no background class. With that in mind, we tried modifying the loss of the Affinity-Net, decreasing the background CAMs, and removing the background completely. We measured the precision and the recall of the predictions. The results showed that removing the background altogether gives almost the best precision, but with a far better recall (cf Table. I).

TABLE I: Quality of Segmentation Labels produced by the AffinityNet. Parameters of eq. 5 are $a = 6, b = 2, c = 3$. $\alpha = \infty$ means background has not be taken into account.

α	Precision	Recall
4	85.22	60.40
16	86.52	66.47
32	86.85	67.28
∞	88.38	87.31

Furthermore, we use the best generated labels to train 3 different segmentation networks, and compare the results with a fully-supervised training on the same networks (cf Table. II).

TABLE II: Results of semantic segmentation on the DEEP-GLOBE land classification challenge. Full labels are labels provided for the challenge . Weak labels are labels produced using weakly supervised model.

Method	Score (mIoU) %
PSPNet Full labels	43.99
- Weak labels	42.97
U-Net Full labels	42.44
- Weak labels	39.25
SegNet Full labels	37.10
- Weak labels	37.76

C. Semantic Segmentation

If we follow the method described in [1], we cannot create the Segmentation Labels without providing the image-level labels, even after training. We modified the method of [1] to directly perform the segmentation, without having to train

⁴<https://competitions.codalab.org/competitions/18468>

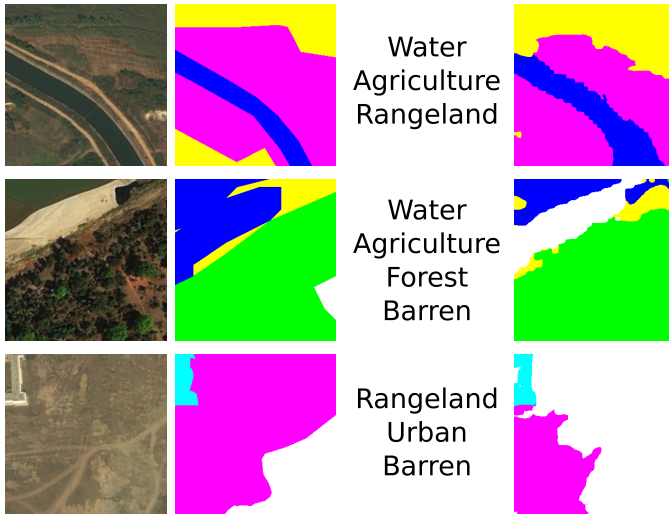


Fig. 3: Example of segmentation labels. From left to right: Original Image, Ground Truth Label, Our Image-Level label, Predicted Label. Note that in order to generate the predicted label, we used only the image-level annotations and original images.

a separate segmentation network. We were able to evaluate our results and compare them to the best, fully-supervised approaches of the deepglobe competition (cf Table. III).

TABLE III: Ranking of our Weakly-Supervised method among Fully-Supervised ones

Method	Score (mIoU)	Rank
Deep Aggregation Net	53.58	1
Dense Fusion Classmate Network	52.64	2
Ours - Weakly Supervised w/o background	45.90	14
Ours - Weakly Supervised w/ background	32.32	33

IV. CONCLUSION

We adapted the method proposed in [1] to our dataset composed of aerial images. The quality of the pixel-level labels generated from image-level labels are quite good. Both offer similar performances when used for training several segmentation network. This means that less information in a dataset does not necessarily implies inferior results.

Furthermore, the semantic segmentation performed by our modification of the AffinityNet gives remarkable results, close to best fully supervised ones.

Further work will be done to improve our results, both on the quality of the generated segmentation labels and the direct semantic segmentation.

Weakly supervised semantic segmentation of satellite images is a cornerstone to the transition from geolocalized text information to complete semantic segmentation.

REFERENCES

- [1] Jiwoon Ahn and Suha Kwak. Learning Pixel-level Semantic Affinity with Image-level Supervision for Weakly Supervised Semantic Segmentation. In *IEEE Comput. Vis. Pattern Recognit.*, 2018.
- [2] Nicolas Audebert, Bertrand Le Saux, and Sébastien Lefèvre. Joint Learning from Earth Observation and OpenStreetMap Data to Get Faster Better Semantic Maps. In *IEEE Conf. Comput. Vis. Pattern Recognit.*, 2017.
- [3] Vijay Badrinarayanan, Alex Kendall, and Roberto Cipolla. SegNet: A Deep Convolutional Encoder-Decoder Architecture for Image Segmentation. *IEEE Trans. Pattern Anal. Mach. Intell.*, 2017.
- [4] Loris Bazzani, Alessandro Bergamo, Dragomir Anguelov, and Lorenzo Torresani. Self-Taught Object Localization with Deep Networks. In *IEEE Winter Conf. Appl. Comput. Vis.*, 2016.
- [5] Amy Bearman, Olga Russakovsky, Vittorio Ferrari, and Li Fei-Fei. What’s the point: Semantic segmentation with point supervision. *Eur. Conf. Comput. Vis.*, 2016.
- [6] Ilke Demir, Krzysztof Koperski, David Lindenbaum, Guan Pang, Jing Huang, Saikat Basu, Forest Hughes, Devis Tuia, and Ramesh Raskar. DeepGlobe 2018: A Challenge to Parse the Earth through Satellite Images. Technical report, 2018.
- [7] Mark Everingham, S. M.Ali Eslami, Luc Van Gool, Christopher K.I. Williams, John Winn, and Andrew Zisserman. The Pascal Visual Object Classes Challenge: A Retrospective. *Int. J. Comput. Vis.*, 2014.
- [8] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep Residual Learning for Image Recognition. In *IEEE Comput. Vis. Pattern Recognit.*, 2016.
- [9] Anna Khoreva, Rodrigo Benenson, Jan Hosang, and Matthias Hein. Simple Does It : Weakly Supervised Instance and Semantic Segmentation. In *IEEE Conf. Comput. Vis. Pattern Recognit.*, 2017.
- [10] Alex Krizhevsky, Ilya Sutskever, and Geoffret E Hinton. ImageNet Classification with Deep Convolutional Neural Networks. In *Adv. Neural Inf. Process. Syst.*, 2012.
- [11] Di Lin, Jifeng Dai, Jiaya Jia, Kaiming He, and Jian Sun. ScribbleSup: Scribble-Supervised Convolutional Networks for Semantic Segmentation. In *IEEE Conf. Comput. Vis. Pattern Recognit.*, 2016.
- [12] TY Lin, Michael Maire, Serge Belongie, and James Hays. Microsoft COCO: Common Objects in Context. In *Eur. Conf. Comput. Vis.*, 2014.
- [13] Diego Marcos, Michele Volpi, Benjamin Kellenberger, and Devis Tuia. Land cover mapping at very high resolution with rotation equivariant CNNs: Towards small yet accurate models. *ISPRS J. Photogramm. Remote Sens.*, 2018.
- [14] Seong Joon Oh, Rodrigo Benenson, Anna Khoreva, Zeynep Akata, Mario Fritz, and Bernt Schiele. Exploiting saliency for object segmentation from image level labels. *IEEE Comput. Vis. Pattern Recognit.*, 2017.
- [15] George Papandreou, Liang-Chieh Chen, Kevin Murphy, and Alan L. Yuille. Weakly- and Semi-Supervised Learning of a DCNN for Semantic Image Segmentation. In *IEEE Int. Conf. Comput. Vis.*, 2015.
- [16] Pedro O. Pinheiro and Ronan Collobert. From image-level to pixel-level labeling with Convolutional Networks. *IEEE Conf. Comput. Vis. Pattern Recognit.*, 07-12-June:1713–1721, 2015.
- [17] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-Net: Convolutional Networks for Biomedical Image Segmentation. In *Int. Conf. Med. Image Comput. Comput. Interv.*, 2015.
- [18] Evan Shelhamer, Jonathan Long, and Trevor Darrell. Fully Convolutional Networks for Semantic Segmentation. *IEEE Trans. Pattern Anal. Mach. Intell.*, 2016.
- [19] Jamie Sherrah. Fully Convolutional Networks for Dense Semantic Labelling of High-Resolution Aerial Imagery. Technical report, 2016.
- [20] Paul Vernaza and Manmohan Chandraker. Learning random-walk label propagation for weakly-supervised semantic segmentation. *IEEE Conf. Comput. Vis. Pattern Recognit.*, 2017.
- [21] Zifeng Wu, Chunhua Shen, and Anton van den Hengel. Wider or Deeper: Revisiting the ResNet Model for Visual Recognition. Technical report, 2016.
- [22] Hengshuang Zhao, Jianping Shi, Xiaojuan Qi, Xiaogang Wang, and Jiaya Jia. Pyramid scene parsing network. *IEEE Conf. Comput. Vis. Pattern Recognit.*, 2017.
- [23] Bolei Zhou, Aditya Khosla, Agata Lapedrizza, Aude Oliva, and Antonio Torralba. Learning Deep Features for Discriminative Localization. In *IEEE Comput. Vis. Pattern Recognit.*, 2016.