



A Cycle GAN Approach for Heterogeneous Domain Adaptation in Land Use Classification

Claire Voreiter, Jean-Christophe Burnel, Pierre Lassalle, Marc Spigai, Romain Hugues, Nicolas Courty

► To cite this version:

Claire Voreiter, Jean-Christophe Burnel, Pierre Lassalle, Marc Spigai, Romain Hugues, et al.. A Cycle GAN Approach for Heterogeneous Domain Adaptation in Land Use Classification. 2020. hal-02537135

HAL Id: hal-02537135

<https://hal.science/hal-02537135>

Preprint submitted on 21 Apr 2020

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.

A CYCLE GAN APPROACH FOR HETEROGENEOUS DOMAIN ADAPTATION IN LAND USE CLASSIFICATION

*Claire Voreiter¹, Jean-Christophe Burnel¹, Pierre Lassalle², Marc Spigai³,
Romain Hugues³, Nicolas Courty¹*

¹Université de Bretagne Sud - IRISA, Rue Yves Mainguy, 56000 Vannes - France

²CNES, 18 avenue Edouard Belin, 31400 Toulouse Cedex 9 - France

³Thales Alenia Space, 26 Avenue Jean François Champollion, 31100 Toulouse - France

ABSTRACT

In the field of remote sensing and more specifically in Earth Observation, new data are available every day, coming from different sensors. Leveraging on those data in classification tasks comes at the price of intense labelling tasks that are not realistic in operational settings. While domain adaptation could be useful to counterbalance this problem, most of the usual methods assume that the data to adapt are comparable (they belong to the same metric space), which is not the case when multiple sensors are at stake. Heterogeneous domain adaptation methods are a particular solution to this problem. We present a novel method to deal with such cases, based on a modified cycleGAN version that incorporates classification losses and a metric space alignment term. We demonstrate its power on a land use classification tasks, with images from both Google Earth and Sentinel-2.

Index Terms— Heterogeneous Domain Adaptation, Cycle GAN, Land Use Classification

1. INTRODUCTION

In remote sensing, a lot of data are available but disposing of associated labels is usually hard and time-consuming. Furthermore, related to the sensor from which the data come from, they do not share the same characteristics. While leveraging on one labelled dataset (the source domain) to perform classification on a new dataset (target) with similar characteristics can be cast as a domain adaptation problem, handling data sources from different sensors is more tricky. This problem is known as an heterogeneous domain adaptation task. In this case, the data have different characteristics (that can include spatial or spectral resolutions) and evolve in different domains. Heterogeneous domain adaptation (HDA) is usually divided in two categories : semi-supervised HDA, when a (small) quantity of target labels are available, and unsupervised HDA, when no target labels at all is available. The difficulty of heterogeneous domain adaptation is to find a sensitive way of linking together the two different domains.

In the state of the art, methods dealing with HDA are mostly shallow methods. EGW and SGW [1] aim to learn an optimal coupling between the two domains through a Gromov-Wasserstein distance. SHFA [2] aims at augmenting source and target samples based on two projection matrices and train simultaneously an SVM classifier on the augmented data. CDLS [3] aims to find representative landmarks to learn a domain-invariant feature subspace and then train a classifier in the learned subspace for target data. DCA [4] aims to jointly obtain a discriminative correlation subspace defined by CCA and then learn a classifier in this subspace. Those methods are not directly amenable to deep learning and the related import quantity of data.

Our new approach for HDA, that is well suited in a deep learning context, is based on the concept of CycleGAN [5, 6, 7]. This method is based on the notion of Generative Adversarial Networks (GAN) [8] which consists in putting two networks in competition: a generator network that produces data and a discriminator network which tries to distinguish the generated sample from a true data. The core idea of CycleGAN is to find an invertible mapping between two domains where examples are provided. Contrary to GANs, no source of noise is considered, two generators (functions that implement mappings to and from the domains are considered, and two discriminators for both domains. We cast the HDA in this framework, but since we want to perform a classification task we need to preserve labels during the image generation. Our contribution focuses on this issue. With the help of this new labeled images, we achieve a classification task on the target domain. A few set of labeled data are used from the target domain and we aim to reduce as much as possible the number of available target labels.

2. PROPOSED METHOD

Two sets of data are used, source data $X_s = \{x_i^s\}_{i=1}^{N_s}$ with \mathbb{P}_s the data distribution, in the domain $S \in \mathbb{R}^{m \times n \times d}$, with $y_s = \{y_i^s\}_{i=1}^{N_s}$ the class labels associated, and, target data $X_t = \{x_i^t\}_{i=1}^{N_t}$ with \mathbb{P}_t the data distribution, in the domain

$T \in \mathbb{R}^{k \times l \times e}$, with n_{y_t} labels $y_t = \{y_i^t\}_{i=1}^{N_t}$. Our method is based on the CycleGAN idea and the notion of GAN. We are going to present this two methods before introducing our improvements.

2.1. GAN

The idea of GANs [8] consist in putting in competition two networks. A generator network G maps a noise and generate samples, and, a discriminator D network tries to distinguish the generated sample from true data samples. The loss function is a minimax game between the generator and the discriminator. We have x the real data, \mathbb{P}_r the data distribution, $\tilde{x} = G(z)$ with z the input of the generator, sampled from some a noise distribution and \mathbb{P}_G the the model distribution. The loss function \mathcal{L}_{GAN} of a GAN, which is minimized over G and maximized over D , reads:

$$\mathcal{L}_{GAN}(G, D) = \mathbb{E}_{x \sim \mathbb{P}_r} [\log(D(x))] + \mathbb{E}_{\tilde{x} \sim \mathbb{P}_G} [\log(1 - D(\tilde{x}))].$$

The CycleGAN method lies in this idea with the addition of a cycle consistency term during the generation.

2.2. Cycle GAN

The Cycle GAN [6] is a method based on GAN used to perform style transfer. The major advantages of this methods are the weakly supervision and the fact that it does not need paired images. The goal of Cycle GAN is to learn mapping function between the two domains S and T . To achieve this, we have two generators, $G_{s2t} : S \rightarrow T$ and $G_{t2s} : T \rightarrow S$, and two discriminators, D_s which aims to distinguish source images from images generated by G_{t2s} and D_t which aims to distinguish source images from images generated by G_{s2t} . The Cycle GAN statement lies in the cycle consistency idea. It is based on the fact that the generated data should be re-transform to their original form: $G_{t2s}(G_{s2t}(x^s)) \approx x^s$ and $G_{s2t}(G_{t2s}(x^t)) \approx x^t$. The cycle consistency loss measures the difference between the original images and the two times generated images. This loss can be written as :

$$\mathcal{L}_{cycle}(G_{s2t}, G_{t2s}) = \mathbb{E}_{x^s \sim \mathbb{P}_s} \|x^s - G_{t2s}(G_{s2t}(x^s))\| + \mathbb{E}_{x^t \sim \mathbb{P}_t} \|x^t - G_{s2t}(G_{t2s}(x^t))\|,$$

where $\|\cdot\|$ is a suitable norm. With the cycle consistency loss, an adversarial loss is applied for both discriminators. The global loss is then :

$$\mathcal{L}_{cycleGAN}(G_{s2t}, G_{t2s}, D_s, D_t) = \mathcal{L}_{GAN}(G_{s2t}, D_t) + \mathcal{L}_{GAN}(G_{t2s}, D_s) + \lambda \mathcal{L}_{cycle}(G_{s2t}, G_{t2s}).$$

2.3. Proposed architecture

The Cycle GAN method can be apply for our case to generate source images in the target domain or the other way around.

However, because we are handling a classification problem, specific cares have to be taken to ensure a correct transfer of the labels.

2.3.1. Metric alignment, classification and total losses

With a Cycle GAN architecture, two samples can be arbitrarily mapped in the target space and it is difficult to ensure that two nearby samples in the target domain will share a correct label. Yet, most classifiers are using this regularity to enforce generalization. This issue is linked to the notion of regularity of the mapping G , which is for example described in [9]. We add a metric alignment loss that will help us to assure that two close samples from the source domain must be close in the target domain after the generation, and the other way around. The corresponding loss simply reads:

$$\mathcal{L}_{metric}(G) = \mathbb{E}_{(x_i, x_j) \sim \mathbb{P}_s} (d(x_i, x_j) - d(G(x_i), G(x_j)))^2,$$

where d is a distance between the samples in the corresponding domains. We notably simply choose a simple Euclidean distance, but this choice could be improved in the future.

This classification consistency is also ensured by a specific loss based on the output of the classifiers C :

$$\mathcal{L}_{classif}(C, G) = \mathbb{E}_{(x, y)} l(C(G(x)), y),$$

evaluated over labelled samples from both source and target domains (if any), and where $l(\cdot, \cdot)$ is a classification loss (typically cross-entropy in a deep learning setting) that forces the sample coming from the generators to be classified according the true associated labels.

Finally, the global loss of our method is:

$$\begin{aligned} \mathcal{L}(G_{s2t}, G_{t2s}, D_s, D_t) = & \mathcal{L}_{cycleGAN}(G_{s2t}, G_{t2s}, D_s, D_t) + \mathcal{L}_{classif}(C_s) \\ & + \mathcal{L}_{classif}(C_t) + \mathcal{L}_{metric}(G_{s2t}) + \mathcal{L}_{metric}(G_{t2s}). \end{aligned}$$

2.3.2. Overall architecture

In our architecture (see Fig. 1), we recover the elements of the Cycle GAN plus two classifiers C_s and C_t . We pretrain C_s with the source data and their labels. Then, we use it to test the relevance of the generated images coming from G_{t2s} with the help of the available target labels. In the same way, we pretrain C_t with the target images and the labels available (when any), and we use it to test the relevance of the generated images coming from G_{s2t} with the help of the source labels. Final classification on the target data is achieved with a third classifier which is trained independently from the architecture of Fig. 1, and discussed subsequently.

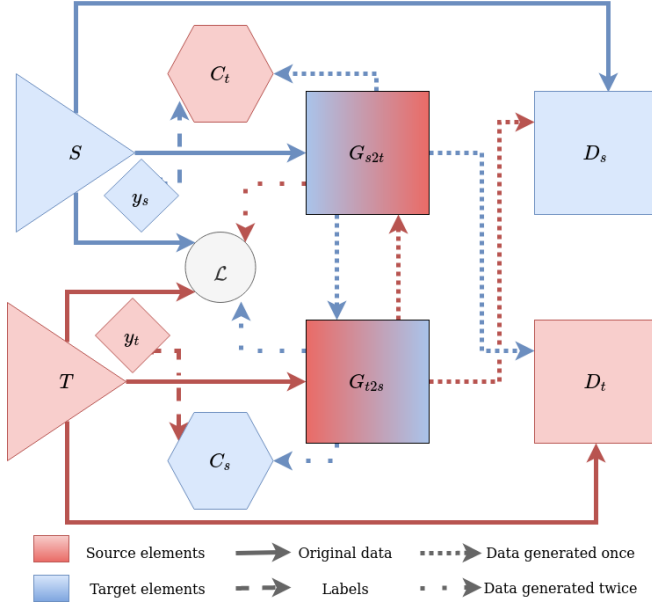


Fig. 1. Diagram of our architecture showing the different elements of its composition.

2.3.3. Final classification

Once the transformation architecture has been trained, there are three ways of performing classification of the unlabelled target samples. In all cases, we use the labelled target data (if any), and we add: *i*) source images transferred to the target domain by G_{s2t} and their corresponding labels. In the remainder, we refer to this method as *HDA_{source}*; *ii*) unlabelled target images, associated to the label obtained by transforming them into their source equivalent (via G_{t2s}) and associated to a label obtained by C_s (*HDA_{target}*); *iii*) a combination of both approaches to generate learning data (*HDA_{full}*). Fig. 3 depicts the three approaches for training the final classifier.

3. EXPERIMENTS AND RESULTS

To assess the qualities of the proposed method, we chose two datasets to make Land Use classification. We selected NWPU-RESISC45 and EuroSAT datasets.

3.1. Datasets

NWPU-RESISC45 [10] contains 31,500 images and cover 45 scene classes. The images come from Google Earth, have a size of 256*256 and are in RGB. The spatial resolution of this dataset varies from about 30 m to 0.2 m per pixel. EuroSAT [11], [12] contains 27,000 images and cover 10 scene classes. The images come from Sentinel-2 satellite, are composed by 13 spectral bands and have a size of 64*64. Since EuroSAT dataset contains only 10 classes, we choose to start

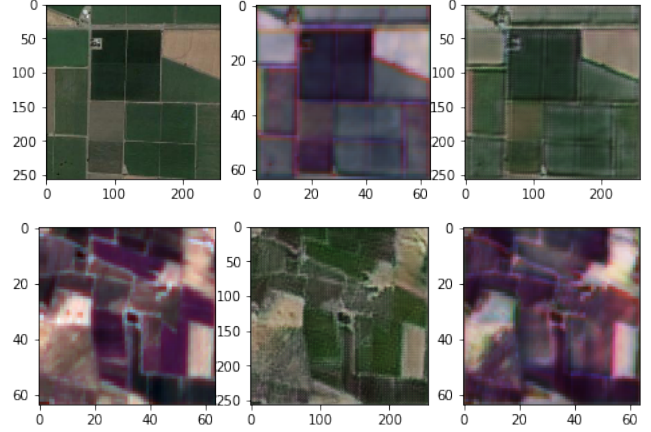


Fig. 2. On the left, we have the original images, in the center, we have the first generated images, and on the right, we have the second generated images which should look like the original images.

from this dataset, we look for common classes between the two datasets. We have 9 classes (annual crop, permanent crop, forest, highway, industrial, pasture, residential, river and sea lake) and we choose to merge annual crop and permanent crop into crop because the images were very similar. Then we selected the corresponding one from NWPU-RESISC45 dataset. For each classes, 700 labelled samples are available. We take 650 samples at best to test our architecture and we keep 50 samples to do the validation.

3.2. Results

We show in Fig. 2 images generated by our method. We want to analyse the influence of the number n_{y_t} of target labels available, so we run our tests with different amount of target labels.

For the model we use a similar architecture as SRGAN [13] for the generators, a convNet for the discriminators made up of four convolutional layers, with batch normalisation and ReLU activation function and also a convNet for the two classifiers made up of four convolutional layers with dropout, maxpooling and ReLU activation function. The final classifier trained is also a convNet made up of three convolutional layers with ReLU activation function. We train the model for 300.000 iterations and the final classifier on 30 epochs.

We compare the accuracy scores of 4 methods: *baseline*, that is the score when only labelled samples in the target are considered, *HDA_{source}*, *HDA_{target}* and *HDA_{full}*. As one can observe from the results of table 1, our method gives better classification results than with target labeled data only in all cases. Less we have available target labels, more we improve the classification results. Interestingly enough, *HDA_{target}* gives the best score consistently. This can be explained by the following fact: since labels are available in the source do-

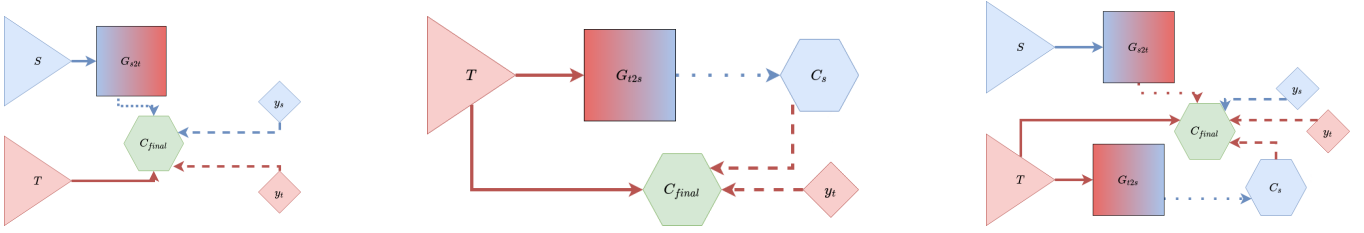


Fig. 3. On the left, diagram of the HDA_{source} method. In the middle, diagram of the HDA_{target} method. On the right, diagram of the HDA_{full} method. These diagrams are showing the origin of the data used for the final classifier.

# n_{y_t}	baseline	HDA _{source}	HDA _{target}	HDA _{full}
650	77.00	71.25	80.00	75.00
325	70.75	70.00	77.00	72.00
130	64.00	50.75	70.00	66.50
65	62.00	43.00	68.50	64.25
1	19.00	32.25	66.75	62.00
0	-	30.75	65.75	57.50

Table 1. Rate of accuracy scores for classification on EuroSAT dataset

main, the C_s classifier is very good, and helps in training a performant G_{ts} transformation. A contrario, the transformation from source to target is less performant, and when combined in the HDA_{full} setting, it only worsens the classification accuracy. Also, the best performance 66% in the fully unsupervised setting is very good (only 11 points behind the golden score (77%), which indicates that our method is capable of transferring knowledge even in the absence of labelled samples in the target domain.

4. CONCLUSIONS

In this paper, we presented a novel approach for classification with heterogeneous domain adaptation using a cycle GAN based approach. Our contribution lies in the addition of specific classification and metric alignment losses, that also helps in the generation process. Experimental results on a Land Use classification problem involving very different remote sensing images indicate the power of our method, in both semi-supervised and unsupervised settings. Future works will consider ablation studies to fully understand the role of each loss independently, and compare with existing unsupervised or semi-supervised deep HDA methods.

5. ACKNOWLEDGEMENT

This work was supported by Centre National d’Études Spatiales (CNES) and Thales Alenia Space.

6. REFERENCES

- [1] Yuguang Yan, Wen Li, Hanrui Wu, Huaqing Min, Minghui Tan, and Qingyao Wu, “Semi-supervised optimal transport for heterogeneous domain adaptation,” in *IJCAI*, 2018, pp. 2969–2975.
- [2] Wen Li, Lixin Duan, Dong Xu, and Ivor W Tsang, “Learning with augmented features for supervised and semi-supervised heterogeneous domain adaptation,” *IEEE transactions on pattern analysis and machine intelligence*, vol. 36, no. 6, pp. 1134–1148, 2013.
- [3] Yao-Hung Hubert Tsai, Yi-Ren Yeh, and Yu-Chiang Frank Wang, “Learning cross-domain landmarks for heterogeneous domain adaptation,” in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2016, pp. 5081–5090.
- [4] Yuguang Yan, Wen Li, Michael KP Ng, Minghui Tan, Hanrui Wu, Huaqing Min, and Qingyao Wu, “Learning discriminative correlation subspace for heterogeneous domain adaptation,” in *IJCAI*, 2017, pp. 3252–3258.
- [5] Anoop Cherian and Alan Sullivan, “Sem-gan: Semantically-consistent image-to-image translation,” in *2019 IEEE Winter Conference on Applications of Computer Vision (WACV)*. IEEE, 2019, pp. 1797–1806.
- [6] Jun-Yan Zhu, Taesung Park, Phillip Isola, and Alexei A Efros, “Unpaired image-to-image translation using cycle-consistent adversarial networks,” in *Proceedings of the IEEE international conference on computer vision*, 2017, pp. 2223–2232.
- [7] Paolo Russo, Fabio M Carlucci, Tatiana Tommasi, and Barbara Caputo, “From source to target and back: symmetric bi-directional adaptive gan,” in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2018, pp. 8099–8108.
- [8] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio, “Generative adversarial nets,” in *Advances in neural information processing systems*, 2014, pp. 2672–2680.
- [9] Emmanuel de Bézenac, Ibrahim Ayed, and Patrick Gallinari, “Optimal unsupervised domain translation,” *CoRR*, vol. abs/1906.01292, 2019.
- [10] Gong Cheng, Junwei Han, and Xiaoqiang Lu, “Remote sensing image scene classification: Benchmark and state of the art,” *Proceedings of the IEEE*, vol. 105, no. 10, pp. 1865–1883, 2017.
- [11] Patrick Helber, Benjamin Bischke, Andreas Dengel, and Damian Borth, “Eurosat: A novel dataset and deep learning benchmark for land use and land cover classification,” *arXiv preprint arXiv:1709.00029*, 2017.
- [12] Patrick Helber, Benjamin Bischke, Andreas Dengel, and Damian Borth, “Introducing eurosat: A novel dataset and deep learning benchmark for land use and land cover classification,” in *IGARSS 2018-2018 IEEE International Geoscience and Remote Sensing Symposium*. IEEE, 2018, pp. 204–207.
- [13] Christian Ledig, Lucas Theis, Ferenc Huszar, Jose Caballero, Andrew Cunningham, Alejandro Acosta, Andrew Aitken, Alykhan Tejani, Johannes Totz, Zehan Wang, et al., “Photo-realistic single image super-resolution using a generative adversarial network,” in *Proceedings of*

the IEEE conference on computer vision and pattern recognition, 2017,
pp. 4681–4690.