



HAL
open science

Experience feedback using Representation Learning for Few-Shot Object Detection on Aerial Images

Pierre Le Jeune, Mustapha Lebbah, Anissa Mokraoui, Hanene Azzag

► **To cite this version:**

Pierre Le Jeune, Mustapha Lebbah, Anissa Mokraoui, Hanene Azzag. Experience feedback using Representation Learning for Few-Shot Object Detection on Aerial Images. 20th International Conference on Machine Learning and Applications (ICMLA) 2021, Dec 2021, Pasadena (virtual), United States. hal-04098829

HAL Id: hal-04098829

<https://hal.science/hal-04098829>

Submitted on 16 May 2023

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.

Experience feedback using Representation Learning for Few-Shot Object Detection on Aerial Images

Pierre Le Jeune
COSE & LIPN & L2TI
Université Sorbonne Paris Nord
pierre.le-jeune@cose.fr

Mustapha Lebbah
LIPN
Université Sorbonne Paris Nord
mustapha.Lebbah@lipn.univ-paris13.fr

Anissa Mokraoui
L2TI
Université Sorbonne Paris Nord
anissa.mokraoui@univ-paris13.fr

Hanene Azzag
LIPN
Université Sorbonne Paris Nord
azzag@lipn.univ-paris13.fr

Abstract—This paper proposes a few-shot method based on Faster R-CNN and representation learning for object detection in aerial images. The two classification branches of Faster R-CNN are replaced by prototypical networks for online adaptation to new classes. These networks produce embeddings vectors for each generated box, which are then compared with class prototypes. The distance between an embedding and a prototype determines the corresponding classification score. The networks are trained in an episodic manner. A new detection task is randomly sampled at each epoch, consisting in detecting only a subset of the classes annotated in the dataset. This strategy encourages the network to adapt to new classes as it would at test time. In addition, several ideas are explored to improve the proposed method such as a hard negative examples mining strategy and self-supervised clustering for background objects. The performance of our method is assessed on DOTA, a large-scale remote sensing images dataset. The experiments conducted provide a broader understanding of the capabilities of representation learning. It highlights in particular some intrinsic weaknesses for the few-shot object detection task. Finally, some suggestions and perspectives are formulated according to these insights.

Index Terms—Faster R-CNN, Few-shot learning, Object detection, Remote sensing images, Representation learning.

I. Introduction

Object detection is a key problem in computer vision. It consists in finding all occurrences of objects belonging to a predefined set of classes in an image and classify them. Its applications range from medical diagnosis to aerial intelligence through autonomous vehicles. Object detection methods automate repetitive and time-consuming tasks performed by human operators until now. In the context of Remote Sensing Images (RSI), detection is used for a wide variety of tasks such as environmental surveillance, urban planning, crops and flock monitoring or traffic analysis.

Deep learning and especially convolutional neural networks (CNNs) outperform previous methods on most computer vision tasks and object detection is no exception. Plenty of methods have been introduced to address this challenge. Among them, Faster R-CNN [1] and YOLO [2] may be the most well-known and studied. These methods perform well when provided with sufficient annotated

data. Yet, this condition is often unmet in practice and the creation of large dataset for detection requires both time and expertise preventing the deployment of such methods for many use cases. Another limitation to the widespread deployment of detection techniques is the lack of adaptability. This is critical for some applications which need to detect different objects from one usage to another. Aerial intelligence is an example of such application: each mission may have its specific objects of interest and therefore a detection model must be adaptable on the fly. The overall objective of this work is to be deployed on vertical aerial images. Yet, large-scale dataset of such images, annotated for object detection, are rare. Few-Shot Learning (FSL) techniques have been introduced to address these issues and deal with limited data. Its principle is to learn general knowledge from a large dataset so that it can generalize efficiently (i.e. quickly and from limited data) on new classes. There exist different approaches for this task, among those transfer learning and meta-learning are the most well-known.

This work introduces a new few-shot learning method for object detection and evaluates its performance on aerial images. It detects objects from only a few examples of a class and without any fine-tuning. The main idea is inspired from prototypical networks [3] which learn an embedding function that maps images into a representation space. The classic Faster R-CNN framework is modified to perform few-shot detection based on this idea. Both classification branch in Region Proposal Network (RPN) and in Fast R-CNN are replaced by prototypical networks to allow fast online adaptation. In addition, a few improvements are introduced on the prototypical baseline in order to fix its weaknesses.

This paper begins with an overview of the scientific literature on object detection, few-shot learning and their intersection. Then, the prototypical Faster R-CNN architecture is presented in detail alongside with several improvements on our baseline. Next, the potential of the proposed modifications throughout a series of experiences

is demonstrated. Finally, the approach is critically discussed, and it is asked whether representation learning is suitable for object detection.

II. Related work

A. Object detection

During the last decade, CNN have made impressive progresses in most computer vision tasks, and object detection is no exception. Among the various methods proposed, two categories emerge. One-stage method with a trade-off on speed and two-stages methods which focus on accuracy. YOLO [2] and Faster R-CNN [1] are the most well-known techniques for each category respectively.

As our work is mainly based on Faster R-CNN, its functioning is described in details. It is made of two stages: a Region Proposal Network (RPN) and a prediction head described in [4]. The RPN is a fully convolutional network, it outputs an objectness score o_a for each anchor a . This score represents the likeliness of having an object within the corresponding patch in the image. In addition, the RPN outputs box shifts b_a^R . The regressions, combined with the anchors sizes and positions, give the actual boxes coordinates in the image. Then, the best scoring boxes are selected to be fed to the prediction head. This head computes refined coordinates b_j^H and classes scores c_j for each box. Following this, a post-processing step filters out small, low-scoring and redundant boxes.

The training of this method is straightforward, each network has two losses, one for the regression branch and one for the classification as described below:

$$\mathcal{L}_{reg}^R(b_i^R, \hat{b}_i^R) = \text{SmoothL1Loss}(b_i^R, \hat{b}_i^R), \quad (1)$$

$$\mathcal{L}_{obj}^R(o_i, \hat{o}_i) = \hat{o}_i \log(o_i) + (1 - \hat{o}_i) \log(1 - o_i), \quad (2)$$

$$\mathcal{L}_{reg}^H(b_j^H, \hat{b}_j^H) = \text{SmoothL1Loss}(b_j^H, \hat{b}_j^H), \quad (3)$$

$$\mathcal{L}_{cls}^H(c_j, \hat{c}_j) = -\log(c_j), \quad (4)$$

where SmoothL1loss is a slight modification to L1 loss function so that the network is less penalized from small regression errors. Hats denote ground truth values.

B. Few-shot learning

Few-shot learning corresponds to learning a task in a limited data setting. Specifically, a task is defined as K -shots, N -ways learning when the training set only contains K examples for each of its N classes. In FSL literature, it is common to introduce the query and support sets for a given task. The support contains the available examples: K images for each of the N classes of the task.

There exists different techniques to tackle low data regime. Transfer learning is one of them. It consists in training a network on a large-scale dataset (source domain) and then fine-tuning it on the few examples (target domain) available for the actual task. These methods require re-training each time a new class is added. In the case of aerial surveillance, this is not suitable as the adaptability must be almost immediate.

Meta-learning is another one. It can be understood as learning the learning process. Two models, a teacher and a student, are learned at the same time. The teacher helps the student to learn quicker and perform well on different tasks with few data available. For instance, the teacher network can be trained to directly output weight updates of the student as described in [5]. Another approach is to output only initial weights for the student as in [6]. While these techniques are promising, they do not scale very well for large networks as the teacher must be substantially larger than the learner.

Another drawback is that most meta-learning methods require a fine-tuning step to be deployed on a specific task. Some exceptions are based on representation learning (or metric learning). Prototypical networks [3] is a pioneer work in using this for FSL. The principle is to train a network to output an abstract representation from an input image. Before inference, a prototype for each class is computed by the embedding network, from the support images. During inference, the query images are embedded and the distance between their representations and the prototypes determine the classification scores. In [3], this is done with a linear classifier, but other choices can be made. Relation networks [7] proposes another network to compute the class scores from the image representation and the prototypes. These methods are usually trained by randomly sampling tasks at each epoch, just as other meta-learning methods. This succession of new tasks helps the network to generalize well and improves its accuracy on unseen classes.

C. Few-shot object detection

The previous section focused only on classification tasks. Detection is a harder problem and so is FSL. That explains why the combination of both was only studied recently. One early work on this problem is Low-Shot Transfer Detector [8]. It leverages transfer learning in order to refine a pre-trained detection network on a small dataset. A regularization loss is introduced to prevent forgetting the base classes. Similarly, [9] proposes to first pre-train a Faster R-CNN on a base dataset and then fine-tune only the last classification and regression layers with the new classes.

Even if RepMet [10] focuses mostly on few-shot classification, authors have shown that their method can also be applied for detection. Their approach is mainly based on metric learning. During base training, they learn alongside the network's weights a set of representatives for different classes. Classes scores are computed from the distance between an embedding and the representatives as in [3].

Most recent work focuses on meta-learning in order to solve Few-Shot Object Detection (FSOD). For instance, reference [11] trains a one stage detector along with a meta features extractor. This extractor computes a reweighting vector for each class from the support set. When a query image is passed through the detector,




Fig. 1: Architectural view of prototypical Faster R-CNN. Embedding vectors for each of the possible box locations is generated by the RPN. These are compared with class prototypes to output objectness scores. The same process occurs in the second stage to produce classes scores.

the feature maps are channel-wise multiplied with the reweighting vectors to produce class-specific maps. These maps are then passed to the detector’s head. Each map is responsible for detecting objects of the corresponding class. A similar class-attention mechanism, but inside the RPN network is proposed in [12]. In the second stage they use multi-relation heads that combine support set information and query image features in multiple ways inside the classification branch. Different methods for combining support information with the query features exist, e.g. [13] creates a graph between the support vectors and the query embeddings. Then a recurrent layer processes it in order to provide attention between regions of interest and support set images. Alike this, [14] processes a graph, whose final nodes are the reweighting vectors, with a graph convolution network.

With regard to FSOD on remote sensing images, very few methods exist. To the best of our knowledge, only two works have been published on this topic. Reference [15] improves on [11] by adding multi-scale features and predictions. It is crucial for tackling RSI as object size can vary greatly. The second one, as described in [13], makes use of a two-stages detector and compute reweighting vectors with a GRU relation module. Both of the methods provide benchmark on VHR-10 dataset [16], yet the number of novel classes and the different base/novel classes splits are different, making the comparison difficult.

III. Proposed method

To deal with FSOD, prototypical Faster R-CNN is introduced. It is a modified version of Faster R-CNN based on metric learning. The key idea is to replace the classification branches in both stages by prototypical networks. This is related to RepMet [10] that learns representatives only in the second stage. Yet it fixes one major flaw of RepMet, the RPN. Once trained, the RPN specializes on classes seen during training. This means

that objects from new classes are filtered out by the RPN, preventing the second stage to detect them. It has a low recall on unseen classes. This is especially not desirable in FSL. Instead, the RPN should be able to adapt to new classes. Our method is an attempt to fill this shortcoming.

First, a few notations are introduced. Let $\mathcal{C} = [1, N]$ be the set of all classes. In the case of n -ways k -shots learning, each task consists in detecting objects among n classes with k examples per class. At each episode a new detection task i is randomly sampled: $C_i \subset \mathcal{C}$ with $|C_i| = n$. Then, data are sampled from the whole dataset to form a support and a query sets: $S_i = \{(x_1, t_1), \dots, (x_{nk}, t_{nk})\}$ and $Q_i = \{(x_1, t_1), \dots, (x_{nk'}, t_{nk'})\}$. For detection, each image x_j comes with a set of annotations t_j which contains the location and label of all objects in the image. Annotations that do not belong to the episode classes are discarded. To build the support set, for each class $c \in C_i$, we select images containing objects of class c and disregard all other objects (i.e. their annotations are not included in the support set but the image is not masked, so they are still visible). If there are more than one object c in the image, only one is selected.

A. Prototypical Faster R-CNN

We propose to change the output dimension of the classification branches in both the RPN and the head. That way, instead of producing a classification (or objectness) score per box, these networks output embedding vectors. Each vector represents the information contained inside the corresponding box. The computation of the representations is straightforward. Features at several scales are extracted by the backbone (denoted f). These features are then fed to the RPN that computes representation vectors for each location in the feature maps (i.e. each corresponding to an anchor). A 2-layers CNN (shared with the regression branch) is applied on each feature maps, then a RoI Align operation extracts same-size features

for each anchor. Finally, a 2-layers MLP maps these features into the embedding space. The dimension of the representation is $r = 128$ and is kept fixed in all our experiments. We call the RPN embedding pipeline g such that $(f \circ g)(x_j) = z_j$ is the set of all anchors’ embeddings for image x_j .

For the second stage, the best scoring boxes produced by the RPN are selected. Their corresponding features are cropped with RoI Align and fed to a 4-layers MLP that outputs embedding vectors as well. Note that the 2 first layers of the head’s MLP are shared with the regression branch. As for the RPN, let h be the encoding function of the head: $(f \circ h)(x_j) = \tilde{z}_j$. Fig. 1 illustrates the whole representation pipeline and how the network computes the different scores.

Classes prototypes are computed from the support images, with the same network, except that only the example’s box is used for feature pooling. When multiple examples are available for a class, their embeddings are averaged to build one prototype per class: $p_{i,c}$ for the RPN and $\tilde{p}_{i,c}$ for the head.

Assuming that each class is represented by a Gaussian distribution centered on its prototype and given the embeddings of an image and the prototypes, the likelihood for each class is computed as follows:

$$p(x_{j,a}|y_{j,a} = c) = \exp\left(\frac{-d(z_{j,a}, p_c)^2}{2\sigma^2}\right), \quad (5)$$

where $x_{j,a}$ refers to the crop of image x_j with anchor a and $y_{j,a}$ is its corresponding label ($y_{j,a} \in C_i \cup \{\emptyset\}$, with \emptyset representing the background class). d is a distance measure over the representation space. In our experiments, d is the Euclidean distance. $\sigma = 0.5$ is the variance of the distribution and is fixed. This likelihood computation is the same for the RPN and the head. However, in the head, the likelihood of the background class is also computed:

$$p(x_{j,a}|\emptyset) = 1 - \max_{c \in C_i} p(x_{j,a}|c). \quad (6)$$

From this, we derive the objectness score in the RPN and the classification (including background) scores in the head:

$$o_{j,a} = \max_{c \in C_i} p(x_{j,a}|c), \quad p(c|x_{j,a}) = \frac{p(x_{j,a}|c)}{\sum_{c \in C_i \cup \{\emptyset\}} p(x_{j,a}|c)}. \quad (7)$$

The training is done episodically, sampling a random subset of classes $C_i \subset \mathcal{C}$ at each epoch. The embedding network is trained using the same loss functions as Faster R-CNN (see (1)–(4)). The loss is computed on the query set and between each update of the network, the prototypes are recomputed from the same support set. Once all query images have been seen by the network, a new task is sampled. In our experiments, the query set contains 5 images for each of its n classes. The optimization is done with Adam optimizer and a learning rate of $1e - 4$. The

backbone network is pretrained on ImageNet and its first layers are kept frozen during training.

B. Iterative improvements

In order to improve the performance of our model, a few improvements are described on top of the baseline.

1) Hard example mining: One issue encountered with the baseline was the detection of all training classes regardless of support examples. This is class memorization. Although this improves performance for training classes, it produces lots of false positive detections. In order to address this, we propose to sample hard negative examples to encourage the network to detect support classes only. The main idea is to take advantage of the annotations for classes not selected in the current task to find hard negative examples, i.e. classes that the network could have memorized from previous tasks. With a new task at each epoch, it is likely that the network still produces detection for objects annotated in one of the previous epochs if it does not rely on the support information. Therefore, these annotations can be used to find examples that should be considered as background for the current task only, in contrast with background examples that do not contain any class of the dataset (i.e. easy negative examples).

2) Moving average prototypes: Another issue with the baseline is that the prototypes can change abruptly, either when the network is updated or when the support set changes. This makes the training unstable. In order to prevent such rapid modification of the prototypes, an exponential moving average is introduced to smooth the modification. Hence, $\bar{p}_c^{(t+1)} = \alpha p_c + (1 - \alpha)\bar{p}_c^{(t)}$ with $\alpha = 0.1$ in our experiments. $\bar{p}_c^{(t)}$ is the averaged prototype for class c at iteration t , while $p_c^{(t)}$ is the prototype computed from the support set, for class c at iteration t (as described in section III-A).

3) Background clustering: Lastly, the baseline shows poor separation of unseen classes embeddings. This leads to poor performance on novel classes at test time. In order to solve this, an inspiration is drawn from [17]. At each iteration, they fit a K-means on the learned representations. This gives pseudo-labels to train the network for classification in a self-supervised manner. Similarly, we propose to fit a K-means on the negative embeddings (i.e. representing boxes not matched by any ground truth object). From the resulting pseudo-labels a contrastive loss function is computed.

IV. Results and experiments

In order to assess the performance of our method, one dataset is chosen: DOTA [18]. It has 16 different classes and contains 400k annotated objects distributed in 2800 RSI.

The experimental protocol is as follows: three classes are reserved for evaluation and two different splits were randomly selected. The network is trained episodically with the remaining classes for 30k iterations. Training

TABLE I: Mean average precision over 5 runs on DOTA dataset with 95% confidence interval. Results are given for two different train/test classes split. Split A: {0, 1, 4}, Split B: {7, 11, 13} (only test classes are given).

k	Split A		Split B	
	Train	Test	Train	Test
1	0.275 \pm 0.01	0.047 \pm 0.02	0.415 \pm 0.03	0.08 \pm 0.01
3	0.352 \pm 0.02	0.024 \pm 0.01	0.392 \pm 0.03	0.101 \pm 0.02
5	0.390 \pm 0.01	0.038 \pm 0.01	0.434 \pm 0.02	0.121 \pm 0.01
10	0.384 \pm 0.02	0.041 \pm 0.01	0.414 \pm 0.03	0.101 \pm 0.02

more improves the performance on training classes, but the network starts to overfit and performs far worse on test classes. Hence, early stopping, to preserve generalization on new classes. The networks are evaluated on both the base and novel classes to assess both the learning and the generalization capabilities. For each experiment, mean average precision (mAP) is provided, computed according to PASCAL VOC [19], with different number of shots: 1, 3, 5 and 10.

A. Results on DOTA

Table I contains the mean average precision reported on DOTA dataset with our best model, according to section IV-B. The results are reported as the mean over 5 runs with 95% confidence intervals. We choose two different train/test classes split randomly. From this, it can be seen that performance on test classes is far below the train classes. This is expected since no supervision is available during training for these objects. It is interesting to see that the more examples (i.e. shots) are given to the network for a class, the better it performs. This may be explained as more shots provide a better approximation of the cluster’s center. Yet this pattern is not always observed, for instance with training classes of split B, the performance is stable with respect to the number of shots. This could happen when a single class is represented by two separate clusters. Having multiple representations for one class increase the chance of sampling representation from different clusters. Hence, the average can be outside both clusters. Nevertheless, it can be seen from Fig. 2 that the prototypes are often not located at the center of their cluster and averaging is often a right strategy to improve performance.

It can also be seen that the increase of performance stagnates with the number of shots. Important gain is reported between 1, 3 and 5 shots but no significant improvement from 5 to 10. A relatively low number of examples is able to approximate correctly the class prototype. Increasing this number does not improve the positioning of the centers any more and can even be harmful as discussed above.

More broadly, the performance is quite low both for base and novel classes and this does not meet our objectives. In comparison, Faster R-CNN trained in a supervised manner on the complete dataset achieved around 0.7 mAP. It

would be unfair to directly compare this value with the performance of our network as it was mainly designed for adaptability. Yet the performance loss on training classes is quite large.

No other method proposed benchmark on DOTA dataset for FSOD and the very few works [13], [15] in this field did not provide their code, thus comparison was not possible.

B. Ablation study

In order to validate the hypothesis formulated in III-B, an ablation study is conducted. The results of this analysis can be found in Table II. They only partially validate this hypothesis. On the one hand, the introduction of hard examples mining and moving average prototypes improves consistently the test mAP in the one-shot setting. On the other hand, background clustering reduces greatly the performance on train classes, while achieving similar results on test classes. According to this analysis, we choose to fix our architecture with hard example mining and the moving average as it combines best train and test performance. Yet, those conclusions must be taken carefully since the variability between different runs is high and the performance gains are small.

TABLE II: Ablation study on improvements described in section III-B. Each row corresponds to the addition on top of the baseline. HEM corresponds to hard example mining, MA to moving average prototypes and BC to background clustering. Once again results are provided for different numbers of shots, both for train and test classes. Bold scores correspond to highest mAP for each k , either for train or test classes.

	k shots	1	3	5	10
	Baseline	Train	0.355	0.359	0.343
Test		0.021	0.027	0.038	0.041
+ HEM	Train	0.312	0.356	0.412	0.343
	Test	0.04	0.023	0.033	0.026
+ HEM + MA	Train	0.265	0.339	0.37	0.351
	Test	0.069	0.035	0.042	0.059
+ HEM + MA + BC	Train	0.133	0.145	0.182	0.148
	Test	0.043	0.041	0.047	0.026

V. Discussions and perspective

From these results, one question arises: is representation learning a suitable choice for object detection? Representation learning methods are competitive with state-of-the-art for few-shot classification, but seem to be inappropriate for few-shot detection. This may be because detection task requires distinguishing between more closely related images. When a trained RPN classifies two overlapping patches of an image, it may produce completely different outputs whether it contains an object entirely or not. Yet, the two patches are visually similar and must be close in the embedding space. This implies a small margin across the RPN decision boundary. Fig. 2a, where threadlike structures can be seen, illustrates

this well. These patterns are made of embeddings from close overlapping patches within the same input image. This prior spatial organization at initialization may be the cause of the low performance of our method. For few-shot classification, there is no such prior organizing the representation space as two different images cannot belong to the same larger image. Hence, the structures and colors are not as similar as those of two close patches for object detection. Fig. 2 shows that training is able to overcome this and organizes the space into semantic clusters. Yet this only happens for training classes, for which strong supervision is available during training. For test classes, the weak supervision available is not enough to build a semantically-aware structure: the test classes representations are mixed together and with negative examples representations (in black in Fig. 2).

Results provided in section IV-A are computed on a test set, whose images were not seen during training. Yet, mAP both on base and novel classes is tracked during training on images already seen by the network. It shows a large improvement compared to what is reported in table I (around 0.65 mAP on train classes and 0.2 on test classes). This strong overfitting showcases the lack of generalization of our method. It may be explained by a simple reason. Objects only represent a small part of the patch for object detection. Hence, much information is embedded in the representation alongside with the relevant semantic information. The position of the embedding is partially controlled by the background and the network can easily learn correlation between background and semantic, making it easy to detect objects correctly in previously seen images. This is especially true for small objects. In Fig. 2b, classes 1, 9 and 10 represent small vehicles and have the largest clusters. Their embeddings have much more variance than the other classes and therefore detection is harder. In comparison, classes 5, 6 and 14 respectively represent basketball court, running track and soccer field. Their clusters are much tighter. Of course, this could also occur for classification but in practice the background variety is far smaller and represents a smaller portion of the image. Hence, it produces tighter clusters and better class separation. It can also be seen from Table I that the performance on split A (containing small vehicles) is worse than on split B. Such a large difference between the splits suggests that the evaluation protocol is not well-suited for this problem. More splits should be used in order to assess generalization on all classes and performance metrics should be reported per class.

In order to leverage strong supervision, one could try self-supervised methods. It has recently been shown that these methods, e.g. [17] can learn generic representations that generalize for many visual tasks, in particular in low-data regime. It would be interesting to investigate further these methods for few-shot object detection, this is planned as future work. In addition, we plan to try methods based on attention mechanism instead of rep-

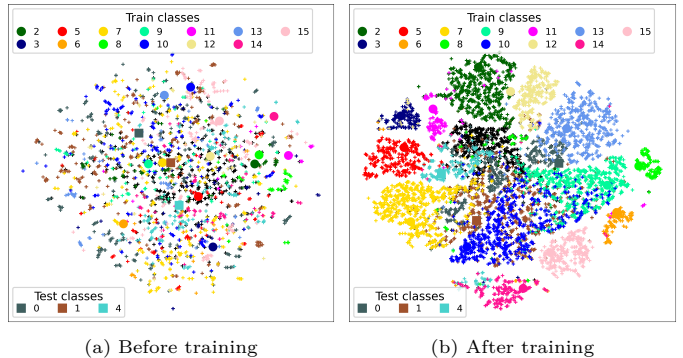


Fig. 2: TSNE visualization on the embedding space, before and after training. Training organizes this space semantically and reduces the threadlike patterns representing close patches in the input image. Yet this is not completely solved for unseen classes during training.

resentation learning as our experiments highlight some weaknesses of the latter. Attention mechanisms, have recently shown great performance for plenty of tasks including few-shot object detection. Finally, a change of the underlying detection architecture is required as modifications in Faster R-CNN can be cumbersome (due to its two stages and the generation of anchors boxes). Instead, FCOS [20], which is a one-stage and anchor-less detector, is probably better suited.

VI. Conclusion

In a nutshell, a novel method is proposed for few-shot object detection based on representation learning. These early results do not meet our expectations in terms of performance. Yet the insights generated in this study allow to understand the strength and weaknesses of representation learning for few-shot object detection task. This will be helpful for future research. Ongoing work is focusing on improving these results using simpler architecture like FCOS.

Acknowledgment

The authors would like to thank COSE for their close collaboration and the funding of this project.

References

- [1] Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun. Faster r-cnn: Towards real-time object detection with region proposal networks. *Advances in neural information processing systems*, 28:91–99, 2015.
- [2] Joseph Redmon, Santosh Divvala, Ross Girshick, and Ali Farhadi. You only look once: Unified, real-time object detection. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 779–788, 2016.
- [3] Jake Snell, Kevin Swersky, and Richard Zemel. Prototypical networks for few-shot learning. In I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 30. Curran Associates, Inc., 2017.
- [4] Ross Girshick. Fast r-cnn. In *Proceedings of the IEEE international conference on computer vision*, pages 1440–1448, 2015.

- [5] S Ravi H Larochelle. Optimization as a model for few-shot learning. In 5th International Conference on Learning Representations, ICLR 2017, Toulon, France, April 24-26, 2017, Conference Track Proceedings, 2017.
- [6] Chelsea Finn, Pieter Abbeel, and Sergey Levine. Model-agnostic meta-learning for fast adaptation of deep networks. In International Conference on Machine Learning, pages 1126–1135. PMLR, 2017.
- [7] Flood Sung, Yongxin Yang, Li Zhang, Tao Xiang, Philip HS Torr, and Timothy M Hospedales. Learning to compare: Relation network for few-shot learning. In Proceedings of the IEEE conference on computer vision and pattern recognition, pages 1199–1208, 2018.
- [8] Hao Chen, Yali Wang, Guoyou Wang, and Yu Qiao. Lstd: A low-shot transfer detector for object detection. In Proceedings of the AAAI Conference on Artificial Intelligence, volume 32, 2018.
- [9] Xin Wang, Thomas E. Huang, Trevor Darrell, Joseph E Gonzalez, and Fisher Yu. Frustratingly simple few-shot object detection. July 2020.
- [10] Leonid Karlinsky, Joseph Shtok, Sivan Harary, Eli Schwartz, Amit Aides, Rogerio Feris, Raja Giryes, and Alex M Bronstein. Repmet: Representative-based metric learning for classification and few-shot object detection. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pages 5197–5206, 2019.
- [11] Bingyi Kang, Zhuang Liu, Xin Wang, Fisher Yu, Jiashi Feng, and Trevor Darrell. Few-shot object detection via feature reweighting. In Proceedings of the IEEE/CVF International Conference on Computer Vision, pages 8420–8429, 2019.
- [12] Qi Fan, Wei Zhuo, Chi-Keung Tang, and Yu-Wing Tai. Few-shot object detection with attention-rpn and multi-relation detector. In CVPR, 2020.
- [13] Zixuan Xiao, Jiahao Qi, Wei Xue, and P. Zhong. Few-shot object detection with self-adaptive attention network for remote sensing images. *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing*, 14:4854–4865, 2021.
- [14] Geonuk Kim, Hong-Gyu Jung, and Seong-Whan Lee. Few-shot object detection via knowledge transfer. In 2020 IEEE International Conference on Systems, Man, and Cybernetics (SMC), pages 3564–3569.
- [15] Xiang Li, Jingyu Deng, and Yi Fang. Few-shot object detection on remote sensing images. *IEEE Transactions on Geoscience and Remote Sensing*, pages 1–14, 2021.
- [16] Gong Cheng, Peicheng Zhou, and Junwei Han. Learning rotation-invariant convolutional neural networks for object detection in vhr optical remote sensing images. *IEEE Transactions on Geoscience and Remote Sensing*, 54(12):7405–7415, 2016.
- [17] Mathilde Caron, Piotr Bojanowski, Armand Joulin, and Matthijs Douze. Deep clustering for unsupervised learning of visual features. In Proceedings of the European Conference on Computer Vision (ECCV), pages 132–149, 2018.
- [18] Gui-Song Xia, Xiang Bai, Jian Ding, Zhen Zhu, Serge Belongie, Jiebo Luo, Mihai Datcu, Marcello Pelillo, and Liangpei Zhang. Dota: A large-scale dataset for object detection in aerial images. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pages 3974–3983, 2018.
- [19] Mark Everingham, Luc Van Gool, Christopher KI Williams, John Winn, and Andrew Zisserman. The pascal visual object classes (voc) challenge. *International journal of computer vision*, 88(2):303–338, 2010.
- [20] Zhi Tian, Chunhua Shen, Hao Chen, and Tong He. Fcos: Fully convolutional one-stage object detection. In Proceedings of the IEEE/CVF International Conference on Computer Vision, pages 9627–9636, 2019.