

Adaptive Self-Training for Object Detection

Renaud Vandeghen
University of Liège
r.vandeghen@uliege.be

Gilles Louppe
University of Liège
g.louppe@uliege.be

Marc Van Droogenbroeck
University of Liège
m.vandroogenbroeck@uliege.be

Abstract

Deep learning has emerged as an effective solution for solving the task of object detection in images but at the cost of requiring large labeled datasets. To mitigate this cost, semi-supervised object detection methods, which consist in leveraging abundant unlabeled data, have been proposed and have already shown impressive results. These methods however often rely on a thresholding mechanism to allocate pseudo-labels. This threshold value is usually determined empirically for a dataset, which is time consuming and requires a new and costly parameter search when the domain changes. In this work, we introduce a new teacher-student method, named Adaptive Self-Training for Object Detection (ASTOD), which is simple and effective. ASTOD selects pseudo-labels adaptively by examining the score histogram. In addition, we also introduce the idea to systematically refine the student, after training, with the labeled data only to improve its performance. While the teacher and the student of ASTOD are trained separately, in the end, the refined student replaces the teacher in an iterative fashion.

Our experiments show that, on the MS-COCO dataset, our method consistently outperforms other adaptive state-of-the-art methods, and performs equally with respect to methods that require a manual parameter sweep search, and are therefore of limited use in practice. Additional experiments with respect to a supervised baseline on the DIOR dataset containing satellite images lead to similar conclusions, and prove that it is possible to adapt the score threshold automatically in self-training, regardless of the data distribution. The code is available at <https://github.com/rvandeghen/ASTOD>.

1. Introduction

On the path to consolidate on the successes of supervised deep learning on large labeled datasets, semi-supervised learning is growing in interest to leverage unlabeled data and improve the performance in many computer vision areas, when the amount of labeled data is scarce. Particularly, semi-supervised learning has led to many improvements for

the task of image classification [2, 3, 18, 28, 29, 33], and is currently growing in interest for object detection. According to current state-of-the-art research [7, 10, 17, 25, 27, 34], semi-supervised learning methods for object detection (SSOD) are usually based on the principle of self-training, wherein a teacher model is first trained with the labeled data in order to generate pseudo-labels for unlabeled data. Then a second model, called the student, is trained with the pseudo-labeled data. Most of the time, the teacher and the student are trained at the same time in a mutual way.

How can we effectively endorse candidate labels generated by methods in the context of SSOD? This question becomes particularly important when considering state-of-the-art classification methods applied to object detection tasks. More precisely, one has to answer the question of how far endogenous (candidate) labels created by a teacher are to be trusted so that, when added to the labeled dataset, the detection performance of a student network twinned with the teacher network can be improved. Keeping only trustable labels by thresholding the predictions provided by the teacher based on their confidence scores is a simple yet effective method. But beyond this simplicity, determining the adequate threshold value remains tricky. Current works in SSOD often require a costly parameter sweep across different values to determine a suitable threshold. While it is easy to understand the behavior of such a threshold regarding the generation of false positives or false negatives, it is not clear which threshold to choose, as evidenced by previous works where the reported optimal threshold value ranges between 0.5 and 0.9 depending on the datasets and network architectures. Also, most works only cover the case of natural scenes, such as MS-COCO [15] and PASCAL VOC [4], preventing drawing conclusions for other datasets.

In this paper, we introduce our Adaptive Self-Training for Object Detection method (ASTOD) to perform the task of object detection. The main idea of our method is to determine the threshold value applied to the confidence score to select pseudo-labels adaptively which is based on the score histogram of the pseudo-labels. In addition, this strategy has the benefit of determining a threshold value for each class without additional cost, which would be very costly

with a parameter sweep for most practical semi-supervised setups. On top of this adaptive threshold, we use different views of the unlabeled images during the generation of pseudo-labels to improve the predictions of the teacher by reducing the number of missed objects, and to improve the predictions of the bounding box coordinates. It is also important to account for the uncertainty in the pseudo-labeled data when we use them. To do so, we downscale the contribution of pseudo-labels in the loss based on their confidence scores. Lastly, we refine our student with the labeled dataset before using it as our new teacher in an iterative way.

In Section 3, we delve into the details of our ASTOD method, whose pipeline is illustrated in Figure 1, after a formal definition of our problem statement. Later, in Section 4, we validate the principle of adaptive self-training with ASTOD for two experimental setups, namely *COCO-standard* and *DIOR*.

Our contributions can be summarized as follows.

- We present a novel end-to-end SSOD method, called ASTOD, based on an iterative teacher-student framework. This method includes a computational-free heuristic based on the score histogram to determine the threshold value for the selection of pseudo-labels.
- We show that using multiple views to generate candidate labels is a simple yet effective technique to improve the labeling process.
- We show that the systematic use of a refinement step is crucial to improve the performance of the student.
- We demonstrate its effectiveness against state-of-the-art methods for two setups.

2. Related Work

Semi-supervised learning. Semi-supervised learning has already been thoroughly studied for image classification. Among the achievements, some methods are based on the principles of consistency training [1, 2, 3, 8, 28, 33], which forces the invariance of a model with respect to input noise by introducing a regularization loss for the unlabeled data. For example, Zhai *et al.* [32] used consistency training to improve the robustness of the model under adversarial attacks. Xie *et al.* [28] have tried another approach in which they minimize the divergence between the output prediction of an unlabeled image and its augmented counterpart.

Another principle for semi-supervised learning is self training [9, 12, 19, 21, 23, 29]. It consists of three parts. First, a teacher model is trained with the labeled data. Then, the trained teacher model is used to generate pseudo-labels on the unlabeled data. Finally, a student model is trained with a dataset comprising the original labels and the pseudo-labels. In particular, Xie *et al.* [29] showed that adding noise

during the training of the student model and increasing the network capacity lead to state-of-the-art results.

Semi-supervised object detection. Driven by the successes obtained for image classification, different semi-supervised learning methods have been tailored for the specific task of the object detection [6, 7, 12, 16, 17, 10, 23, 25, 27, 34, 36], even though pioneering work began in 2005 [21]. Among them, most methods [7, 10, 12, 16, 17, 23, 27, 34] use a threshold value determined empirically to select or reject a pseudo-label. Only few of them are designed without threshold [6, 25]. Particularly, one of the first work was done by Jeong *et al.* [6], who proposed a consistency-based semi-supervised learning method by applying consistency between a horizontally flipped image and the original one for the classification part, with the Jensen-Shannon divergence, as well as for the localization part, with a weighted sum of the squared errors of the four different components of the localization loss. They applied this consistency loss for labeled and unlabeled data. Given that this loss can be dominated by the background class, they performed a background elimination, which removes predictions likely to belong to the background. Another approach is to use soft pseudo-labels [25], which means that the whole distribution of class probabilities is used rather than the hard pseudo-label. Those methods give more flexibility as they do not need any threshold value. Among the threshold-based methods, the field of SSOD has grown in interest after that Sohn *et al.* [23] introduced a semi-supervised learning method based on self-training and augmentation-driven consistency regularization. They start by training a teacher in a supervised manner. Afterwards, they use the teacher to generate candidate labels, which are selected when their confidence scores are above a threshold of 0.9. The method then uses a second model, which is trained on both the labeled and pseudo-labeled data by jointly minimizing a conventional supervised loss and a weighted unsupervised loss based on consistency regularization with strong data augmentations. In [12], Li *et al.* presented a selective self-supervised self-training for object detection method. They started with a teacher-student self-training method with a threshold value of 0.7 during pseudo-labeling, and improved the pseudo-labeling step with a so-called selective network. This network splits the set of pseudo-labels into three categories (positive, negative, and ambiguity), but only the positive pseudo-labels are considered in the loss term. They also implemented a consistency term in their loss based on the work presented in [6], which is also only used for the positive pseudo-labels. Liu *et al.* [16] proposed an unbiased teacher, which is a teacher-student method trained in a mutual setting. The teacher generates pseudo-labels for the training of the student and, then, the teacher is updated with exponential moving average (EMA), leading to continually improving models. The

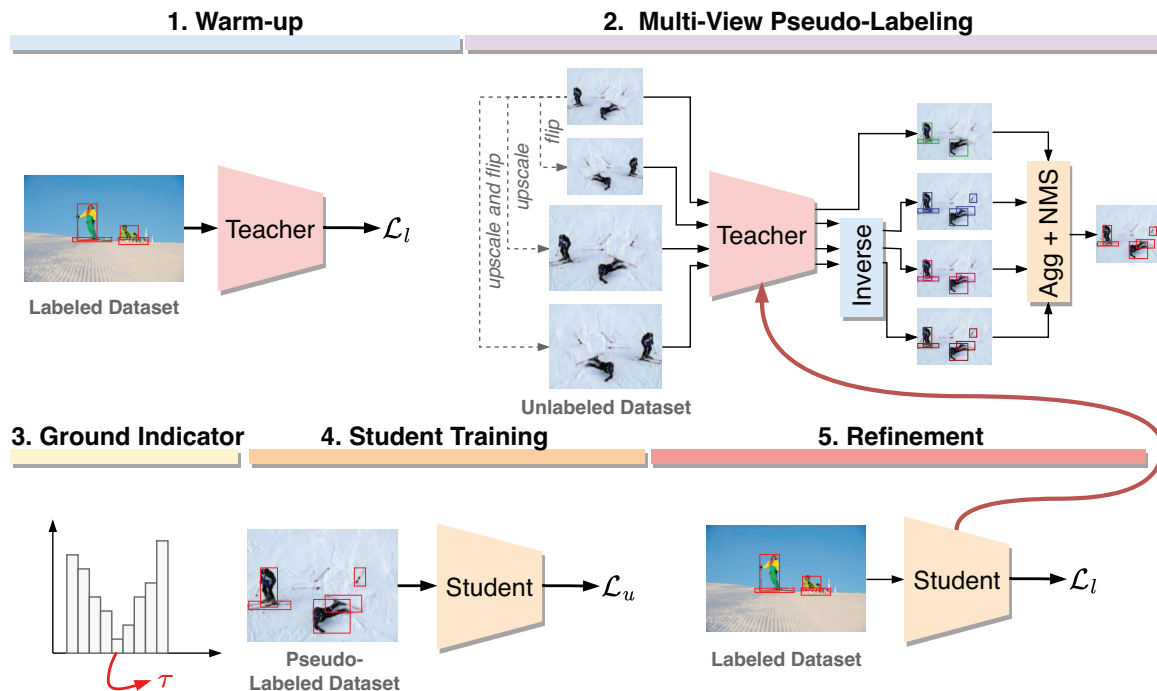


Figure 1: Pipeline of our self-training ASTOD method. (1) A teacher is trained with the labeled dataset. (2) We use the teacher to generate candidate labels on the unlabeled data using multiple views. We apply the inverse view transformation to gather the different predictions in the same dimensional space. The predictions are then merged with NMS. (3) Based on the confidence score histogram, we determine the threshold value τ to filter the candidate boxes, leading to a pseudo-labeled dataset. (4) Next, we train the student with the labeled and pseudo-labeled datasets. (5) Finally, we refine the student with the labeled dataset and use it to replace the teacher. ASTOD can then be used in an iterative fashion by replacing the teacher (2) with the refined student.

authors also used a threshold value of 0.7 to remove boxes with low confidence scores and they addressed the problem of class imbalance by replacing the standard cross-entropy loss with the focal loss [14]. They also published a second version of their unbiased teacher [17] for anchor-free detectors. Zhang *et al.* [34] also addressed the problem of class imbalance with two modules. The first module addresses the problem of foreground-background imbalance by pasting synthetic objects from the training/pseudo dataset in the training images. The second module addresses the problem of foreground-foreground imbalance, which changes the sampling probability with respect to the class occurrence. Kim *et al.* [7] presented a data augmentation technique for the unlabeled dataset that mixes image tiles and feature tiles together and then unmixes the features for the student. Those unmixed feature maps are then processed by the RPN and ROI heads with the pseudo-labels being generated by the teacher with the same weakly augmented images. Li *et al.* [10] also adopted the teacher-student dual learning but took into account the noisy nature of pseudo-boxes regression. Their method is based on a learning

scheme that uses multiple views for both the images and the feature maps to enforce consistency. Tanaka *et al.* [24] recently proposed to optimize the threshold based on the β -score and without iterating on the student.

The current literature on the topic of semi-supervised learning for object detection exhibits a wide variety of nuances around a teacher-student scheme and the calculation of pseudo-labels. Alongside, this often results in the heuristic determination of parameters among which the determination of a threshold on the confidence score during the generation of pseudo-labels. As opposed to most approaches of current literature, we intentionally skip this process thanks to an adaptive calculation of such threshold embedded into a new iterative, multiple-view, and refined teacher-student scheme. This forms the basis of our concept of adaptive self-training. In the case of ASTOD, this calculation occurs by analyzing the histogram of scores associated to the generation of pseudo-labels in a fully automatic fashion.

3. Method

Problem statement. We consider a set, \mathcal{D} , of images, \mathbf{x}_i , containing several classes of objects to be detected. Among \mathcal{D} , only a subset of images, \mathbf{x}_i^l , are annotated with the class and localization of all objects of interest, \mathbf{y}_i^l (called “labels” or ground-truths), and compose the subset of labeled images $\mathcal{D}_l = \{\mathbf{x}_i^l, \mathbf{y}_i^l\}_{i=1}^{N_l}$, where N_l is the number of labeled images in \mathcal{D} . Each ground-truth \mathbf{y}_i^l is composed by a set of classes c and bounding box coordinates \mathbf{b} . The remaining images of \mathcal{D} with no labels, \mathbf{x}_i^u , compose the subset of unlabeled images $\mathcal{D}_u = \{\mathbf{x}_i^u\}_{i=1}^{N_u}$, where N_u is the number of unlabeled images in \mathcal{D} . These subsets are complementary sets (that is $\mathcal{D} = \mathcal{D}_l \cup \mathcal{D}_u$), and we assume that they come from the same data distribution. In semi-supervised learning setups, we often have $N_l \ll N_u$.

Teacher warm up. Our method relies on a teacher-student scheme, where the student learns from the pseudo-labels generated by the teacher. Thus, the first step is to learn a teacher that is able to generate high-quality candidate labels, which are all the predictions made by a model without restrictions. The first step of our method then consists in training the teacher model with the labeled data only. We use the conventional training loss for object detection, which is the sum of the classification and regression losses:

$$\mathcal{L}_l = \sum_{i=1} \left[\sum_{j=i} (\mathcal{L}_{cls}(p(c_j|\mathbf{x}_i), c_j) + \mathcal{L}_{reg}(p(\mathbf{b}_j|\mathbf{x}_i), \mathbf{b}_j)) \right], \quad (1)$$

where the index j corresponds to the index of an anchor, $p(c_j|\mathbf{x}_i)$ is the predicted class probability of anchor j in the image \mathbf{x}_i , and $p(\mathbf{b}_j|\mathbf{x}_i)$ are the 4 bounding box coordinates of a predicted bounding box.

Multi-view pseudo-labeling and ground threshold. After we warm up the teacher model, we use it to generate candidate labels for the unlabeled data. Since we want to mitigate the false negatives due to missed predictions and we want the most accurate predictions, the inference of each unlabeled image is processed under multiple views: original image, horizontally flipped image, rescaled image, and both flipped and rescaled image. Afterwards, we apply the inverse transformations to the predictions so that they can be aggregated in the same dimensional space. To reduce redundancy, we apply non-maximum suppression (NMS) on each prediction before and after aggregation. This leads to a subset of candidate pseudo-labels $\mathcal{D}_{\hat{u}} = \{\mathbf{x}_i^u, \hat{\mathbf{y}}_i^u\}_{i=1}^{N_c}$, with N_c being the number of unlabeled images which have at least one prediction $\hat{\mathbf{y}}$, given that we automatically discard images without prediction. Note that each prediction $\hat{\mathbf{y}}_i^u$ is composed by the set of predicted classes, its corresponding bounding box coordinates, and the confidence scores \mathbf{s} associated to each box. Now that we have access to high-quality candidate labels, we need to select among

them those that can be considered as true positives. In contrast to classification tasks, where we can select the class with the highest probability, this is a challenging step for SSOD. Indeed, there can be multiple objects in the same image, meaning that an independent decision must be taken for each anchor. The most straightforward and, by far, most common solution is to threshold the candidates based on their score predictions.

Previous works in the field use a threshold value, denoted by τ , which suits at best their method and the dataset on which they evaluate it. Typical values for this threshold range between 0.5 and 0.9. However, this threshold value is often determined with a costly parameter sweep, unique for all classes, and is optimized for only one image distribution (natural scenes with MS-COCO). To account for those shortcomings, we propose a new heuristic, called *ground thresholding*, based on the score histogram to determine the threshold value: ground thresholding selects the bin with the lowest density. From our experience, taking the bin with the lowest density is a heuristic that constitutes a well-suited compromise solution for eliminating false positives (bins on the left) while preserving a high enough recall (bins on the right). The final pseudo-labeled dataset is then $\mathcal{D}_p = \{\mathbf{x}_i^p, \hat{\mathbf{y}}_i^p\}_{i=1}^{N_p}$, with N_p being the number of candidate images which have at least one prediction that satisfies $\hat{\mathbf{y}}^c \geq \tau$. Since this heuristic does not require the burden of a parameter sweep to find the threshold value, it can be applied independently for each class, which does not bias the threshold value with respect to the class occurrence. It also means that our method can easily generalize to any dataset without any additional computational cost.

Iterative student training. We train the student model in the same manner as for the teacher, but with the labeled and pseudo-labeled data ($\mathcal{D}' = \mathcal{D}_l \cup \mathcal{D}_p$). During the training of the student, we do not distinguish images coming from \mathcal{D}_l or \mathcal{D}_p . However, to account for the uncertainty in the pseudo-labeled data, we generalize the weighting term

$$\alpha_j = \begin{cases} \frac{s_j - \tau_l}{\tau_h - \tau_l} & \text{if } \tau_l \leq s_j < \tau_h, \\ 1 & \text{otherwise.} \end{cases} \quad (2)$$

used in [26] for the loss by fixing $\tau_h = 1$, where τ_h and τ_l represent a high and a low threshold value and s_j the score prediction. This leads to the weighted loss function:

$$\mathcal{L}_u = \sum_{i=1} \left[\sum_{j=i} \alpha_j (\mathcal{L}_{cls} + \mathcal{L}_{reg}) \right], \quad (3)$$

where \mathcal{L}_{cls} and \mathcal{L}_{reg} are the same classification and regression losses as in Equation (1). The weighting factor α_j used to reduce the contribution of each prediction is then defined as

$$\alpha_j = \begin{cases} \frac{s_j - \tau_j}{1 - \tau_j} & \text{if } \tau_j \leq s_j \leq 1, \\ 1 & \text{otherwise,} \end{cases} \quad (4)$$

with s_j and τ_j being the score and the class-wise threshold value associated to the prediction. Since the score value of labeled data are implicitly set to 1, only pseudo-labeled data contribute to the weighting factor of the loss.

The final step consists in the refinement of the trained student model with the labeled data only. Our method can then be used in an iterative pipeline, where the refined student model will become the new teacher. Since we expect that the student model achieves better results compared to the teacher, its predictions for the candidate labels should be of higher quality which thus leading to an even better new student.

4. Experiments

4.1. Experimental setup

Datasets. Our experimental setup follows the methodology introduced in STAC [23]. In particular, we evaluate our method on two setups: (*setup 1*) natural images on MS-COCO [15] and (*setup 2*) satellite images from the DIOR [11] dataset. For the first setup (called *COCO-standard* hereafter), we randomly sample 1, 2, 5 and 10% labeled training data out of the 118k images available in the `train2017` split and use the remaining ones as unlabeled training data. For the second setup (*DIOR*), we first shuffle all the labeled images in two parts: the training part with 80% and the validation part with the remaining 20%. Then, we sample 10% of the training dataset as labeled data and the remaining 90% as unlabeled data. Unlike most of the other works in semi-supervised for object detection, which use PASCAL VOC [4] as second dataset, we evaluate our method on satellite images with the DIOR dataset to analyze our method for a totally different image distribution. For both *COCO-standard* and *DIOR*, we report the mean and standard deviation of the $AP_{50:95}$ (mAP) over 5 folds.

Implementation details. For a fair comparison with previous works, we use Faster-RCNN [20] with FPN [13] and a ResNet-50 [5] backbone pretrained on ImageNet [22] as object detector. For the teacher warm-up, we train the model for 20k steps of gradient descent with a starting learning rate of 0.08 that decays after 13k and 18k steps by a factor 10. For the generation of pseudo-labels, we use 4 different views of the unlabeled image: (1) normal view, (2) upscale of the original image by a factor of 2, (3) horizontal flip of the original image, and (4) both upscaling and flipping of the original image. From the score histogram, we set the threshold by selecting the bin with the lowest density between 0.5 and 1 with 21 bins —the choice of the $[0.5, 1]$ range was motivated by the need to select only pseudo-labels with enough confidence, while we choose 21 bins because we wanted an odd number of bins and, by experience, the impact of more bins on the threshold value was insignificant. The student models are trained for 180k

steps, with the same learning rate as the teacher, which follows the same decay strategy after 120k and 160k steps. Finally, the student models are refined on the labeled dataset for 10k steps with a learning rate starting at 0.0008 which decays after 6k steps. All the models for *COCO-standard* are trained on 4 GPUs, with a batch size of 16 per GPU. We apply random color and scale jitter as data augmentation. When we train the student models, the batches are formed with 2 labeled and 14 unlabeled images. For *DIOR*, we use 3 scale levels per anchor to better match the different bounding box shapes. The batch size is reduced to 8 per GPU and the student is trained for 90k steps.

4.2. Results

COCO-standard. We compare our model with the state-of-the-art semi-supervised object detection methods on *COCO-standard*, as it is the main benchmark adopted by the SSOD community. We group the different methods according to how their threshold value is set, if any. In particular, we group methods that perform a parameter search to find the optimal threshold value. This kind of methods represents most of previous works in the field, such as STAC [23], Soft Teacher [30] or Unbiased Teacher [16, 17]. The second group is for methods that do not need an empirical search for their threshold, such as CSD [6], Humble Teacher [25] and our method. One could say that the former group are dataset dependent while the latter ignore the dataset distribution. Even though our ASTOD method has a threshold parameter, it is adaptive to the dataset, thus closer to the second group than the first. The results are shown in Table 1, where our results are obtained after 3 iterations of student training plus refinement. While being competitive w.r.t. to the state-of-the-art methods with empirical threshold, like Unbiased Teacher v2 [17] and PseCo [10], our method consistently outperforms methods that do not take into account the dataset distribution [6, 25]. It is important to note that if Unbiased Teacher v2 [17] and PseCo [10] have better performances than ASTOD, they would be more challenging to use in practice on a new dataset, simply because there is no data to fine-tune their thresholds.

DIOR. It is important to design SSOD methods that are usable and effective in many different setups. While previous works in the field have mainly focused on natural scene images with MS-COCO and PASCAL-VOC, we decided to evaluate our method with a different and challenging setup. We targeted the field of satellite images because of the growing interest around it, with the DIOR dataset.

Since there is no baseline for SSOD methods for that dataset, we will compare with the supervised baseline achieved by our teacher model. We generated candidate labels with the Flip+Scale strategy and used a class-wise ground threshold. We also refine the student model to further improve its performance. The results of the different

Table 1: Experimental results on *COCO-standard* for the mAP: we report the mean and standard deviation over 5 randomly sampled dataset. We group the different methods w.r.t. to their thresholding strategy. The methods in the middle of the table use a manual empirical search for the threshold value (these methods are thus intractable when applied on a new unknown domain), while methods in the lower part are fully automatic. The results of Supervised† represents the performance of our teacher, which sets our supervised baseline. The results of our method are obtained after 3 iterations of student with refined models, where we used our ground threshold and the multi-views strategies during the pseudo-labeling step.

	1%	2%	5%	10%
Supervised	9.05 ± 0.16	12.70 ± 0.15	18.47 ± 0.22	23.86 ± 0.81
Supervised†	12.14 ± 0.21	16.67 ± 0.30	23.59 ± 0.20	29.34 ± 0.20
STAC [23]	13.97 ± 0.35	18.25 ± 0.25	24.38 ± 0.12	28.64 ± 0.21
Instant Teaching [35]	18.05 ± 0.15	22.45 ± 0.15	26.75 ± 0.05	30.40 ± 0.05
ISMT [31]	18.88 ± 0.74	22.43 ± 0.56	26.37 ± 0.24	30.53 ± 0.52
Unbiased Teacher [16]	20.75 ± 0.12	24.30 ± 0.07	28.27 ± 0.11	31.50 ± 0.10
Soft Teacher [30]	20.46 ± 0.39	-	30.74 ± 0.08	34.04 ± 0.14
Omni-DETR [27]	18.6	23.2	30.2	34.1
Unbiased Teacher v2 [17]	25.40 ± 0.36	28.37 ± 0.03	31.85 ± 0.09	35.05 ± 0.02
PseCo [10]	22.43 ± 0.36	27.77 ± 0.18	32.50 ± 0.08	36.06 ± 0.24
CSD [6]	10.51 ± 0.06	13.93 ± 0.12	18.63 ± 0.07	22.46 ± 0.08
Humble Teacher [25]	16.96 ± 0.38	21.72 ± 0.24	27.70 ± 0.15	31.61 ± 0.28
Ours (ASTOD)	19.47 ± 0.39	24.85 ± 0.21	30.43 ± 0.50	34.58 ± 0.22

Table 2: Comparison between models trained on *DIOR*. The refined student models are trained with candidates labels generated with the Scale+Flip technique and a class-wise threshold. We report the mean and standard deviation over 5 randomly sampled dataset.

	Supervised	Student	Refined
mAP	47.59 ± 0.36	51.23 ± 0.35	52.89 ± 0.33

models are presented in Table 2. The gain obtained after one iteration shows the effectiveness and robustness of our method towards completely different data distribution, meaning that it can be further used in other applications.

4.3. Ablation study

We study our method w.r.t. to its different components on *COCO-standard* with a labeled dataset size of 10%.

Pseudo-labeling. It is important to rely on high-quality pseudo-labels. To obtain those high-quality pseudo-labels, it is possible to use a high threshold value but at the cost of rejecting potential true positives with lower confidence scores. However, it is not possible to avoid false negatives due to missed predictions. Our multi-view pseudo-labeling strategy can help to reduce their numbers. Figure 2 shows the predicted candidate labels for the different views we consider. We can effectively see that only using the normal view fails to predict some objects in the image, such as the right snowboard, and adding the predictions of the other views solves the problem. Since our aggregation of boxes is

performed with NMS, which is a score-based method, the final boxes are a mix of the different views, leading to the best possible candidates.

Ground threshold. The key component of our method is its ability to determine a suited threshold value without empirical search. The proposed strategy is to compute the score histogram and set the threshold value to the bin with the lowest number of instances. While the number of bins and the score range are parameters of the proposed method, Figure 3 shows that the shapes of the histograms are the same, that is U-shaped with high density regions for very low and very high scores, and that they do not influence the position of the lowest density bin. Since this heuristic is independent on the data distribution, it can be applied for each class separately, which gives the possibility to have a set of thresholds rather than a single one. However, training with a uniform threshold seems to achieve better results on *COCO-standard*, as can be seen in Table 4. Looking at Figure 4, which shows the score histogram for a single class and for all the classes jointly, we can see that both of them define a U-shape. This is the shape that we expect since a threshold value lower than the ground threshold would lead to more pseudo-labels, with many of them having a high probability to be false positives. Also, if the score threshold was higher than the ground threshold, we would probably create false negatives. But the problem arises for classes that are hard to learn. For those classes, the chosen heuristic can fail. Since most of the predictions for those classes may have a low confidence score, the histogram can be monotonically decreasing, leading to a ground threshold equal

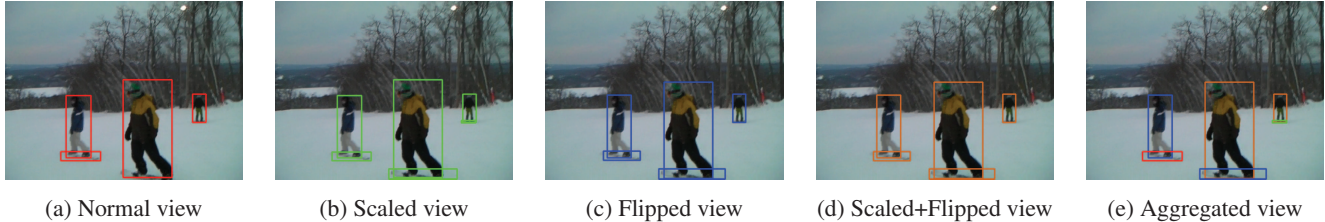
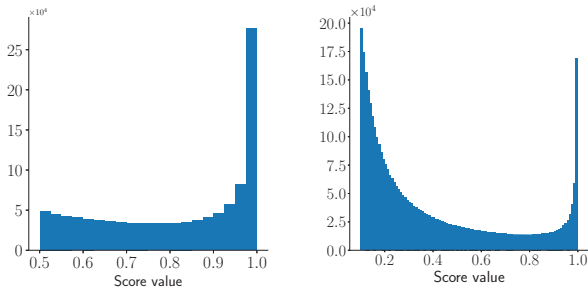


Figure 2: Comparison between the candidate labels for the different views. The normal view (a) misses two snowboards. Both flipped and scaled+flipped views (c) and (d) miss the small snowboard. Only the scaled view (b) has detected all the snowboards. The aggregated view (e) combines the information of all images (with NMS) to produce the final candidate labels. Note that images (b), (c) and (d) are transformed back to the original space.

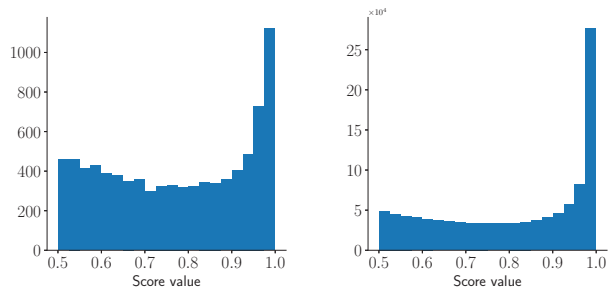


(a) Histogram ranging from 0.5 to 1 with 21 bins. (b) Histogram ranging from 0.1 to 1 with 101 bins.

Figure 3: Histograms for different parameters.

to the last bin. There are multiple consequences to that scenario. First, we generate a lot of false negatives in the pseudo-labeled dataset. Then, the student trains on noisy labels, which will again emphasize the problem of classes hard to learn. Taking a uniform threshold sets the threshold value at a lower score than their ground threshold, leading to fewer false negatives. However, in the *DIOR* setup, we observe that this problem is not present, which we explain by the fact that the teacher is able to better learn the different classes. As shown in Table 4, the student model performs better with a class-wise threshold.

The key advantage of our method is that it eliminates the need for a parameter search to find the threshold value. However, it is interesting to see how it behaves against this parameter search. Table 3 shows the performance of the student model trained with different threshold values and trained with our method. The results depict two interesting behaviors: (1) the optimal threshold from the parameter search does not always gives a better performance, as can be seen on *DIOR*, (2) the optimal threshold value with parameter search for two distinct image distributions does not give the same threshold (0.7 for *COCO-standard* and 0.8 for *DIOR*). This emphasizes that a manual sweep is unsuitable for generalization purposes. It is also important to note that



(a) Score histogram for a single class. (b) Score histogram for all classes.

Figure 4: Score histograms for a single class ($\tau = 0.7$) (a), and for all the classes ($\tau = 0.75$) (b).

we can process our candidate labels using 6 bins ranging from 0.5 to 1. In this setup, the width of a bin is 0.1, meaning that the bin with the lowest density will match a value that could have been selected with this classical grid search. On *COCO-standard*, we observed that the ground threshold value in this particular setup is 0.7, which appears to be the optimal value in Table 4. This result further consolidate that our ground threshold strategy gives a good threshold value at no cost. Although this particular setup would give the best result for the *COCO-standard* setup, we argue that restraining our method to this particular example does not fulfill our idea to be adaptive to any dataset, which is confirmed with the result obtained on *DIOR*.

Table 3: Performances obtained with a parameter search on the threshold compared to our method. We report the mean over 5 randomly sampled dataset for both setups.

τ	0.5	0.6	0.7	0.8	0.9	Ours
<i>COCO</i> 10%	32.81	32.83	33.02	32.82	32.57	32.91
<i>DIOR</i>	52.14	52.29	52.55	52.61	52.49	52.89

Iterative students + Refinement. Since our method is not

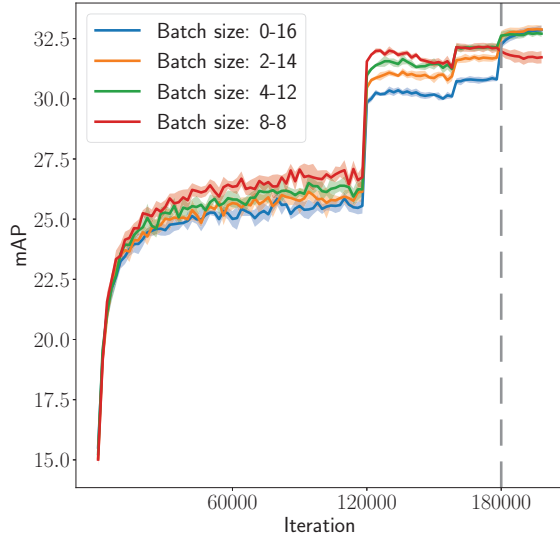


Figure 5: Comparison between the different learning curves of student and refined models w.r.t. the batch size configuration. The vertical dashed line indicate when the refinement step begins.

Table 4: Comparison between refined student models trained with different thresholding strategies. We report the mean and standard deviation over 5 folds for both setups. Surprisingly, using a uniform value performs better than using one determined class-wise on *COCO-standard*.

Setup	Class-wise τ	Uniform τ
<i>COCO-standard</i>	32.56 ± 0.11	32.91 ± 0.16
<i>DIOR</i>	52.89 ± 0.33	52.52 ± 0.25

based on a mutual training between the teacher and student models, the pseudo-labels used for the student are fixed. However, if we can obtain better results by the end of its training, we can expect to have better candidate labels by using the student as our new teacher. Before we replace the student as our new teacher, we refine the student on the labeled dataset only for a few gradient descent steps. This idea has already been used by Vandeghen *et al.* [26] to improve the ROI Heads with only trustable ground-truth labels. As it is shown in Table 5 for *COCO-standard* and in Table 2 for *DIOR*, this final trick is highly effective. The results obtained for 3 iterations, before and after refinement, are shown in Table 5.

During our experiments, we performed an analysis on the batch size distribution between the labeled and pseudo labeled images. The different configurations were 0|16,

Table 5: Comparison between the different iterations of student and refined models for the mAP on *COCO-standard*. There is a twofold message from those results: (1) Consecutive iterations of student training consistently improve the performance compared to the previous iteration. (2) Refining the student is a simple yet effective way to boost the performance. We report the mean and standard deviation over 5 randomly sampled dataset.

	1%	2%	5%	10%
Supervised	9.05 ± 0.16	12.70 ± 0.15	18.47 ± 0.22	23.86 ± 0.81
Supervised†	12.14 ± 0.21	16.67 ± 0.30	23.59 ± 0.20	29.34 ± 0.20
Student 1	16.57 ± 0.46	21.53 ± 0.34	27.64 ± 0.17	31.77 ± 0.14
Refined 1	16.67 ± 0.36	21.93 ± 0.36	28.47 ± 0.43	32.91 ± 0.16
Student 2	17.75 ± 0.31	23.23 ± 0.29	29.17 ± 0.45	32.88 ± 0.18
Refined 2	17.95 ± 0.37	23.62 ± 0.33	29.54 ± 0.45	33.86 ± 0.18
Student 3	18.71 ± 0.30	24.23 ± 0.34	29.65 ± 0.41	33.40 ± 0.23
Refined 3	19.47 ± 0.39	24.85 ± 0.21	30.43 ± 0.50	34.58 ± 0.22

2|14, 4|12 and 8|8, for the labeled and unlabeled size respectively. The averaged learning curves of the first student models are shown in Figure 5, where the training of the student models stops at 180,000 iterations. From the student results, it could be obvious to discard the first two configurations. However, those refined models tend to perform better than the last two configurations. This analysis shows that (1) refining the student models is a crucial step to improve their performance, and (2) drawing some conclusions with only the student performance may not be sufficient.

5. Conclusion

In this paper, we present ASTOD, an iterative end-to-end self-training method for object detection. Our method solves the problem of parameter sweep for the threshold value in SSOD with a heuristic threshold value which adapts easily to different setups. We also present the systematic use of a refinement step of the student models to improve their performance. Our experiments show that our method largely outperforms state-of-the-art methods in SSOD, that are threshold-free methods.

Limitations and further work. While our method shows an excellent capacity to adapt to diverse data distributions, there is still potential to adapt it to methods which approach the teacher-student scheme with mutual learning. We believe that more work should address the problem of thresholding methods based on parameter search. Finally, a deeper investigation regarding the refining step may be useful, as we have shown that this step consistently improves the performance.

Acknowledgments. The present research benefited from computational resources made available on Lucia, the Tier-1 supercomputer of the Walloon Region, infrastructure funded by the Walloon Region under the grant agreement n°1910247.

References

- [1] Philip Bachman, Ouais Alsharif, and Doina Precup. Learning with pseudo-ensembles. In *Adv. Neural Inf. Process. Syst. (NeurIPS)*, volume 27, Montréal, Can., Dec. 2014. Curran Assoc. Inc. [2](#)
- [2] David Berthelot, Nicholas Carlini, Ekin Cubuk, Alex Kurakin, Kihyuk Sohn, Han Zhang, and Colin Raffel. ReMix-Match: Semi-supervised learning with distribution matching and augmentation anchoring. In *Int. Conf. Learn. Represent. (ICLR)*, Addis Ababa, Ethiopia, Apr.-May 2020. [1](#), [2](#)
- [3] David Berthelot, Nicholas Carlini, Ian Goodfellow, Nicolas Papernot, Avital Oliver, and Colin Raffel. MixMatch: A holistic approach to semi-supervised learning. In *Adv. Neural Inf. Process. Syst. (NeurIPS)*, volume 32, Vancouver, Canada, Dec. 2019. Curran Assoc. Inc. [1](#), [2](#)
- [4] Mark Everingham, Luc Van Gool, Christopher K. I. Williams, John Winn, and Andrew Zisserman. The PASCAL visual object classes (VOC) challenge. *Int. J. Comput. Vis.*, 88(2):303–338, Jun. 2010. [1](#), [5](#)
- [5] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *IEEE Int. Conf. Comput. Vis. Pattern Recognit. (CVPR)*, pages 770–778, Las Vegas, NV, USA, Jun. 2016. [5](#)
- [6] Jisoo Jeong, Seungeui Lee, Jeesoo Kim, and Nojun Kwak. Consistency-based semi-supervised learning for object detection. In *Adv. Neural Inf. Process. Syst. (NeurIPS)*, volume 32, Vancouver, Canada, Dec. 2019. Curran Assoc. Inc. [2](#), [5](#), [6](#)
- [7] JongMok Kim, JooYoung Jang, Seunghyeon Seo, Jisoo Jeong, Jongkeun Na, and Nojun Kwak. MUM: Mix image tiles and UnMix feature tiles for semi-supervised object detection. In *IEEE Int. Conf. Comput. Vis. Pattern Recognit. (CVPR)*, pages 14492–14501, New Orleans, LA, USA, Jun. 2022. Inst. Electr. Electron. Eng. (IEEE). [1](#), [2](#), [3](#)
- [8] Samuli Laine and Timo Aila. Temporal ensembling for semi-supervised learning. *ArXiv*, abs/1610.02242, 2016. [2](#)
- [9] Dong-hyun Lee. Pseudo-label: The simple and efficient semi-supervised learning method for deep neural networks. In *ICML Workshop on Challenges in Representation Learning*, volume 3, pages 1–6, Atlanta, Georgia, USA, Jun. 2013. [2](#)
- [10] Gang Li, Xiang Li, Yujie Wang, Yichao Wu, Ding Liang, and Shanshan Zhang. PseCo: Pseudo labeling and consistency training for semi-supervised object detection. *ArXiv*, abs/2203.16317, 2022. [1](#), [2](#), [3](#), [5](#), [6](#)
- [11] Ke Li, Gang Wan, Gong Cheng, Liqiu Meng, and Junwei Han. Object detection in optical remote sensing images: A survey and a new benchmark. *ISPRS J. Photogramm. Remote Sens.*, 159:296–307, Jan. 2020. [5](#)
- [12] Yandong Li, Di Huang, Danfeng Qin, Liqiang Wang, and Boqing Gong. Improving object detection with selective self-supervised self-training. In *Eur. Conf. Comput. Vis. (ECCV)*, volume 12374 of *Lect. Notes Comput. Sci.*, pages 589–607. Springer, Oct. 2020. [2](#)
- [13] Tsung-Yi Lin, Piotr Dollar, Ross Girshick, Kaiming He, Bharath Hariharan, and Serge Belongie. Feature pyramid networks for object detection. In *IEEE Int. Conf. Comput. Vis. Pattern Recognit. (CVPR)*, pages 2117–2125, Honolulu, HI, USA, Jul. 2017. [5](#)
- [14] Tsung-Yi Lin, Priya Goyal, Ross Girshick, Kaiming He, and Piotr Dollar. Focal loss for dense object detection. *ArXiv*, abs/1708.02002, 2017. [3](#)
- [15] Tsung-Yi. Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C. Lawrence Zitnick. Microsoft COCO: Common Objects in Context. In *Eur. Conf. Comput. Vis. (ECCV)*, volume 8693 of *Lect. Notes Comput. Sci.*, pages 740–755. Springer, Sept. 2014. [1](#), [5](#)
- [16] Yen-Cheng Liu, Chih-Yao Ma, Zijian He, Chia-Wen Kuo, Kan Chen, Pzizhao Zhang, Bichen Wu, Zsolt Kira, and Peter Vajda. Unbiased teacher for semi-supervised object detection. In *Int. Conf. Learn. Represent. (ICLR)*, May 2021. [2](#), [5](#), [6](#)
- [17] Yen-Cheng Liu, Chih-Yao Ma, and Zsolt Kira. Unbiased teacher v2: Semi-supervised object detection for anchor-free and anchor-based detectors. In *IEEE Int. Conf. Comput. Vis. Pattern Recognit. (CVPR)*, pages 9809–9818, New Orleans, LA, USA, Jun. 2022. Inst. Electr. Electron. Eng. (IEEE). [1](#), [2](#), [3](#), [5](#), [6](#)
- [18] Hieu Pham, Zihang Dai, Qizhe Xie, Minh-Thang Luong, and Quoc Le. Meta pseudo labels. In *IEEE Int. Conf. Comput. Vis. Pattern Recognit. (CVPR)*, pages 11557–11568, Nashville, TN, USA, Jun. 2021. [1](#)
- [19] Hieu Pham, Zihang Dai, Qizhe Xie, Minh-Thang Luong, and Quoc V. Le. Meta pseudo labels. *ArXiv*, abs/2003.10580, 2020. [2](#)
- [20] Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun. Faster R-CNN: Towards real-time object detection with region proposal networks. *IEEE Trans. Pattern Anal. Mach. Intell.*, 39(6):1137–1149, Jun. 2017. [5](#)
- [21] Chuck Rosenberg, Martial Hebert, and Henry Schneiderman. Semi-supervised self-training of object detection models. In *IEEE Workshops on Applications of Computer Vision (WACV/MOTION)*, volume 1, pages 29–36, Breckenridge, Colorado, USA, Jan. 2005. Inst. Electr. Electron. Eng. (IEEE). [2](#)
- [22] Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng Huang, Andrej Karpathy, Aditya Khosla, Michael Bernstein, Alexander C. Berg, and Li Fei-Fei. ImageNet large scale visual recognition challenge. *Int. J. Comput. Vis.*, 115(3):211–252, Apr. 2015. [5](#)
- [23] Kihyuk Sohn, Zizhao Zhang, Chun-Liang Li, Han Zhang, Chen-Yu Lee, and Tomas Pfister. A simple semi-supervised learning framework for object detection. *ArXiv*, abs/2005.04757, 2020. [2](#), [5](#), [6](#)
- [24] Yuki Tanaka, Shuhei M. Yoshida, and Makoto Terao. Non-iterative optimization of pseudo-labeling thresholds for training object detection models from multiple datasets. In *IEEE Int. Conf. Image Process. (ICIP)*, pages 1676–1680, Bordeaux, France, Oct. 2022. Inst. Electr. Electron. Eng. (IEEE). [3](#)
- [25] Yihe Tang, Weifeng Chen, Yijun Luo, and Yuting Zhang. Humble teachers teach better students for semi-supervised object detection. In *IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, pages 3131–3140, Nashville, TN, USA, Jun. 2021. Inst. Electr. Electron. Eng. (IEEE). [1](#), [2](#), [5](#), [6](#)

- [26] Renaud Vandeghen, Anthony Cioppa, and Marc Van Droogenbroeck. Semi-supervised training to improve player and ball detection in soccer. In *IEEE Int. Conf. Comput. Vis. Pattern Recognit. Work. (CVPRW), CVsports*, New Orleans, LA, USA, Jun. 2022. 4, 8
- [27] Pei Wang, Zhaowei Cai, Hao Yang, Gurumurthy Swaminathan, Nuno Vasconcelos, Bernt Schiele, and Stefano Soatto. Omni-DETR: Omni-supervised object detection with transformers. In *IEEE Int. Conf. Comput. Vis. Pattern Recognit. (CVPR)*, pages 9357–9366, New Orleans, LA, USA, Jun. 2022. Inst. Electr. Electron. Eng. (IEEE). 1, 2, 6
- [28] Qizhe Xie, Zihang Dai, Eduard Hovy, Minh-Thang Luong, and Quoc V. Le. Unsupervised data augmentation for consistency training. In *Adv. Neural Inf. Process. Syst. (NeurIPS)*, volume 33, pages 6256–6268. Curran Assoc. Inc., Dec. 2020. 1, 2
- [29] Qizhe Xie, Minh-Thang Luong, Eduard Hovy, and Quoc V. Le. Self-training with noisy student improves ImageNet classification. In *IEEE Int. Conf. Comput. Vis. Pattern Recognit. (CVPR)*, pages 10684–10695, Seattle, WA, USA, Jun. 2020. 1, 2
- [30] Mengde Xu, Zheng Zhang, Han Hu, Jianfeng Wang, Lijuan Wang, Fangyun Wei, Xiang Bai, and Zicheng Liu. End-to-end semi-supervised object detection with soft teacher. In *IEEE Int. Conf. Comput. Vis. (ICCV)*, pages 3060–3069, Montréal, Can., Oct. 2021. 5, 6
- [31] Qize Yang, Xihan Wei, Biao Wang, Xian-Sheng Hua, and Lei Zhang. Interactive self-training with mean teachers for semi-supervised object detection. In *IEEE Int. Conf. Comput. Vis. Pattern Recognit. (CVPR)*, pages 5941–5950, Nashville, TN, USA, Jun. 2021. 6
- [32] Runtian Zhai, Tianle Cai, Di He, Chen Dan, Kun He, John Hopcroft, and Liwei Wang. Adversarially robust generalization just requires more unlabeled data. *ArXiv*, abs/1906.00555, 2019. 2
- [33] Xiaohua Zhai, Avital Oliver, Alexander Kolesnikov, and Lucas Beyer. S4L: Self-supervised semi-supervised learning. In *IEEE Int. Conf. Comput. Vis. (ICCV)*, pages 1476–1485, Seoul, South Korea, Oct. 2019. Inst. Electr. Electron. Eng. (IEEE). 1, 2
- [34] Fangyuan Zhang, Tianxiang Pan, and Bin Wang. Semi-supervised object detection with adaptive class-rebalancing self-training. *ArXiv*, abs/2107.05031, 2021. 1, 2, 3
- [35] Qiang Zhou, Chaohui Yu, Zhibin Wang, Qi Qian, and Hao Li. Instant-teaching: An end-to-end semi-supervised object detection framework. In *IEEE Int. Conf. Comput. Vis. Pattern Recognit. (CVPR)*, pages 4079–4088, Nashville, TN, USA, Jun. 2021. Inst. Electr. Electron. Eng. (IEEE). 6
- [36] Barret Zoph, Golnaz Ghiasi, Tsung-Yi Lin, Yin Cui, Hanxiao Liu, Ekin D. Cubuk, and Quoc V. Le. Rethinking pre-training and self-training. In *Adv. Neural Inf. Process. Syst. (NeurIPS)*, volume 33, pages 3833–3845. Curran Assoc. Inc., Dec. 2020. 2