

# PROGRESSIVE TRANSFORMATION LEARNING FOR LEVERAGING VIRTUAL IMAGES IN TRAINING

**Anonymous authors**

Paper under double-blind review

## ABSTRACT

To effectively interrogate UAV-based images for detecting objects of interest, such as humans, it is essential to acquire large-scale UAV-based datasets that include human instances with various poses captured from widely varying viewing angles. As a viable alternative to laborious and costly data curation, we introduce Progressive Transformation Learning (PTL), which gradually augments a training dataset by adding transformed virtual images with enhanced realism. Generally, a virtual2real transformation generator in the conditional GAN framework suffers from quality degradation when a large domain gap exists between real and virtual images. To deal with the domain gap, PTL takes a novel approach that progressively iterates the following three steps: 1) select a subset from a pool of virtual images according to the domain gap, 2) transform the selected virtual images to enhance realism, and 3) add the transformed virtual images to the training set while removing them from the pool. In PTL, accurately quantifying the domain gap is critical. To do that, we theoretically demonstrate that the feature representation space of a given object detector can be modeled as a multivariate Gaussian distribution from which the Mahalanobis distance between a virtual object and the Gaussian distribution of each object category in the representation space can be readily computed. Experiments show that PTL results in a substantial performance increase over the baseline, especially in the small data and the cross-domain regime.

## 1 INTRODUCTION

Training an object detector usually requires a large-scale training image set so that the detector can acquire the ability to detect objects' diverse appearances. This desire for a large-scale training set is bound to be greater for object categories with more diverse appearances, such as the human category whose appearances vary greatly depending on its pose or viewing angles. Moreover, a person's appearance becomes more varied in images captured by an unmanned aerial vehicle (UAV), leading to a wide variety of camera viewing angles compared to ground-based cameras, making the desire for a large-scale training set even greater. In this paper, we aim to satisfy this desire, especially when the availability of UAV-based images to train a human detector is scarce, where this desire is more pressing.

As an intuitive way to expand the training set, one might consider synthesizing virtual images to imitate real-world images by controlling the optical and physical conditions in a virtual environment. Virtual images are particularly useful for UAV-based object detection since abundant object instances can be rendered with varying UAV locations and camera viewing angles along with ground-truth information (e.g., bounding boxes, segmentation masks) that comes free of charge. Therefore, a large-scale virtual image set covering diverse appearances of human subjects that are rarely shown in existing UAV-based object detection benchmarks (Zhu et al., 2021; Barekatin et al., 2017; ICG) can be conveniently acquired by controlling rendering entities and parameters in a virtual environment, such as poses, camera viewing angles, and illumination conditions.

To make virtual images usable for training real-world object detection models, recent works (Hoffman et al., 2018; Mu et al., 2020; Li & Lee, 2021; Qiu et al., 2021) transform virtual images to look realistic. They commonly use the virtual2real generator (Mirza & Osindero, 2014) trained with the conditional GAN framework to transform images in the source domain to have the visual properties

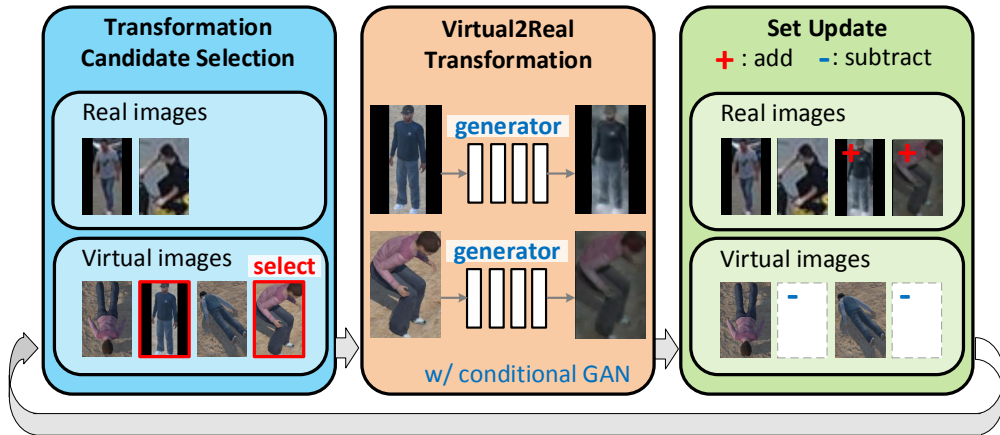


Figure 1: Overview of Progressive Transformation Learning (PTL).

of images in the target domain. Here, virtual and real images are treated as the source and target domains, respectively. However, the large discrepancy in visual appearance between the two domains, referred to as the “domain gap”, result in the degraded transformation quality of the generator. In fact, the aforementioned works using virtual images validate their methods where the domain gap is not large (e.g., digit detection (Hoffman et al., 2018)) or when additional information is available (e.g., animal pose estimation with additional keypoint information (Mu et al., 2020; Li & Lee, 2021; Qiu et al., 2021)). In our case, real and virtual humans in UAV-based images inevitably have a large domain gap due to the wide variety of human appearances.

To address the large domain gap, one critical question inherent in our task is *how one can accurately measure the domain gap*. Consequently, we estimate the domain gap in the representation space of a human detector trained on the real images. The representation space of the detector is learned such that test samples, which have significantly different properties than training samples from the perspective of the detector, are located away from the training samples. In this paper, we show that the feature distribution of object entities belonging to a certain category, such as the human category, in the representation space can be modeled with a multivariate Gaussian distribution if the following two conditions are met: i) the detector uses the sigmoid function to normalize the final output of the detector and ii) the representation space is constructed using the output of the penultimate layer of the detector. This idea is inspired by Lee et al. (2018), which shows that softmax-based classifiers can be modeled as multivariate Gaussian distributions. In this paper, we show that the proposition is also applicable to sigmoid-based classifiers, which are widely used by state-of-the-art (SOTA) detectors. Based on this modeling, when the two aforementioned conditions are met, the human category in the representation space can be represented by two parameters (i.e., mean and covariance) of a multivariate Gaussian distribution that can be computed on the training images. With the empirically calculated mean and covariance, the domain gap from a single virtual human image to real human images (i.e., the training set) can be measured using the Mahalanobis distance (Mahalanobis, 1936).

To add virtual images to the training set to include more diverse appearances of objects while preventing the transformation quality degradation caused by large domain gaps, we introduce *Progressive Transformation Learning (PTL)* (Figure 1). PTL progressively expands the training set by adding virtual images through iterating the three steps: 1) transformation candidate selection, 2) virtual2real transformation, and 3) set update. When selecting transformation candidates from a pool of virtual images, we use weighted random sampling, which gives higher weights to images with smaller domain gaps. The weight takes an exponential term with one hyperparameter controlling the ratio between images with smaller domain gaps and images with more diverse appearances. Then, the virtual2real transformation generator is trained via the conditional GAN framework, taking the selected virtual images (i.e., transformation candidates) as the “source” and the images in the training set as the “target”. After transforming the transformation candidates by applying the virtual2real transformation generator, the training set is expanded with the transformed candidates while the original candidates are excluded from the pool of virtual images.

The main contribution of this paper is that we have validated the utility of virtual images in augmenting training data via PTL coupled with carefully designed comprehensive experiments. We first use the task of low-shot learning, where adequately expanding datasets has notable effects. Specifically, PTL provides better accuracy on three UAV-view human detection benchmarks than other previous works that leverage virtual images in training, as well as methods that only use real images. Then, we validate PTL on the cross-domain detection task where training and test data are from distinct datasets, the problem we frequently face in real-world scenarios. Similarly, PTL consistently presents significantly better accuracy than the baseline. The experimental results indicate that a high-performance human detection model can be effectively learned via PTL, even with the significant lack of training data representing real operational environments.

## 2 RELATED WORKS

**Leveraging virtual images in training.** In this section, we have listed previous works that demonstrate how virtual images can play a role in a variety of real-world applications when used for training. In fact, virtual images are desirable for model training as large-scale labeled datasets can be built virtually almost free of charge. Unfortunately, when virtual images are used for model training without proper care, it is shown that the performance improvement is limited due to the domain gap between the virtual images and the real test images. Generally, previous works leveraging virtual images during model training can be summarized into three approaches according to how they exploit the advantages of virtual images and address the challenges of using virtual images.

The most intuitive and widely used approach of using virtual images is to pre-train a model on virtual images and fine-tune the pre-trained model on real images acquired in the same domain as the test images (Handa et al., 2016; Kim et al., 2016; Gaidon et al., 2016; Fabbri et al., 2018; Varol et al., 2018; Mu et al., 2020; Mishra et al., 2022; Guo et al., 2022; Jin et al., 2022; Baek & Shim, 2022). This approach aims to avoid the domain gap by fine-tuning the model on real images acquired under the same conditions and environments of test images. While the first approach seeks to use the representative capability learned from *large-scale* virtual datasets, the second approach seeks to exploit additional information that can be easily *labeled* on virtual images. For example, Liu et al. (2022) annotates part segmentation maps when acquiring virtual vehicle images, and uses the part segmentation results from the pre-trained model on the virtual images during the fine-tuning process. Similarly, Yan et al. (2022) uses depth and semantic information labeled when acquiring virtual images.

The third approach directly builds training batches consisting of both virtual images and real images. Ros et al. (2016) and Richter et al. (2016) adopt the most naive approach to build a batch by randomly selecting a fixed number of images from each of the real and virtual image sets. However, even if the number of virtual images is several orders of magnitude greater than the number of real training images, this approach does not provide remarkably better accuracy and may even provide worse accuracy than its counterparts using only real images for training. In this case, the effect of using virtual images during model training does not appear as expected because the large domain gap between the real (test) images and the virtual images is not adequately addressed. In this paper, we also use virtual images directly for model training while considering appropriately reducing this domain gap.

**Progressive learning.** Progressive learning is a machine learning strategy that continuously trains a model from easy to hard tasks, primarily for the purpose of training stabilization or fast optimization. One of the most common applications for progressive learning is incrementally increasing the network capacity to improve network capability. The most intuitive approach in this category is to gradually increase the network size (e.g., depth or width) to alleviate the training difficulty of very deep networks (Fahlman & Lebiere, 1989; Bengio et al., 2006; Smith et al., 2016; Chen et al., 2016; Wei et al., 2016; Gong et al., 2019; Li et al., 2022a). Conversely, Zhang & He (2020) uses progressive learning in the direction of reducing the network size for fast training. Progressive learning is also used in GAN frameworks to enhance the generator’s ability to transform input images of larger resolution (Karras et al., 2018). Curriculum learning (Bengio et al., 2009; Spitkovsky et al., 2009; Khan et al., 2011; Graves et al., 2017), which continuously raises the level of training from easy to difficult samples, also falls into this category.

Progressive learning is also used to deal with the scalability of datasets that are incompletely labeled. In a semi-supervised learning task where unlabeled samples can be used in training, Wang et al. (2017), Wang et al. (2022a), and Cao et al. (2022) adopt progressive learning by gradually increasing the number of unlabeled data used for training. Self-learning in Kumar et al. (2010) and Kaul et al. (2022) uses progressive learning by repeating the two steps, assigning labels depending on the current detector and updating the current detector with these labels.

Our method can be seen as similar to the second approach in that it also intend to expand the training dataset. However, our method uses progressive learning to reliably add realistically transformed virtual images to the training dataset by avoiding quality degradation of the transformation, which has never been attempted before using the progressive learning strategy.

### 3 METHOD

#### 3.1 MEASURING THE DOMAIN GAP BETWEEN REAL AND VIRTUAL IMAGES

**Modeling training set with multivariate Gaussian distribution.** The purpose of adopting progressive transformation learning, which progressively expands the training set with a subset of the realistically transformed virtual images instead of expanding it with the full set at once, is to avoid the significant domain gap between the real and virtual images when training the transformation generator. Here, the domain gap is measured in the representation space of the detector, which is learned so that two samples with different properties from the perspective of the detector are separated far from each other.

In general, the representation space of the detector refers to the space formed by the output of the penultimate layer of the detector since all layers of the detector except for the last layer can be transferred for different downstream tasks (He et al., 2020; Chen et al., 2020; Grill et al., 2020). In Lee et al. (2018), it is shown that for the softmax-based classifier, the distribution of each category in the representation space can be modeled as a multivariate Gaussian distribution. Object detector generally uses the sigmoid function (i.e.,  $f_{sigmoid}(\mathbf{x}) = 1/(1 + \exp(-\mathbf{w}_c^T \mathbf{x} - b_c))$  for the category  $c$ ), which does not consider outputs for other categories, instead of the softmax function (i.e.,  $f_{softmax}(\mathbf{x}) = \exp(\mathbf{w}_c^T \mathbf{x} + b_c) / \sum_{c'} \exp(\mathbf{w}_{c'}^T \mathbf{x} + b_{c'})$  for the category  $c$ ) that competes for outputs for all categories to normalize the model output to [0 1]. This is because, unlike classification, the detection task must take into account that two or more co-located objects may be active on a single output. In the supplementary material, we show that even for the sigmoid-based detector, the distribution of each category in the representation space can also be modeled as a multivariate Gaussian distribution.

Specifically, let  $\mathbf{x} \in \mathcal{X}$  and  $y = \{y_c\}_{c=1, \dots, C} \in \mathcal{Y}$ ,  $y_c \in \{0, 1\}$  be an input and its label, respectively. Then, the representation space of the sigmoid-based detector can be expressed as follows:

$$P(f(\mathbf{x})|y_c = 1) \sim \mathcal{N}(f(\mathbf{x})|\mu_c, \Sigma_c), \quad (1)$$

where  $f(\cdot)$  denotes the output of the penultimate layer of the detector.  $\mu_c$  and  $\Sigma_c$  are the mean and the covariance of the multivariate Gaussian distribution for the category  $c$ .  $\mu_c$  and  $\Sigma_c$  can be calculated over the entire set of training images as follows:

$$\mu_c = \frac{1}{|D_c|} \sum_{\mathbf{x} \in D_c} f(\mathbf{x}), \quad \Sigma_c = \frac{1}{|D_c|} \sum_{\mathbf{x} \in D_c} (f(\mathbf{x}) - \mu_c)(f(\mathbf{x}) - \mu_c)^\top, \quad (2)$$

where  $D_c$  is the set of instances for the category  $c$ . Practically, any detection whose IoU with the groundtruth of the category  $c$  is greater than 0.5 belongs to  $D_c$ .

**Measuring domain gap.** After  $\mu_c$  and  $\Sigma_c$  are empirically calculated to represent  $D_c$ , the domain gap between a new image  $\mathbf{x}_{new}$  and  $D_c$  can be measured using the Mahalanobis distance, as follows:

$$d(\mathbf{x}_{new}) = (f(\mathbf{x}_{new}) - \mu_c)^\top \Sigma_c^{-1} (f(\mathbf{x}_{new}) - \mu_c). \quad (3)$$

This measurement of the domain gap is highly dependent on the detector’s ability to detect objects in the image. It is commonly known that the detection capability of a detector is greatly affected by the image size as well as the object appearance in the image. To mitigate the effect of image size

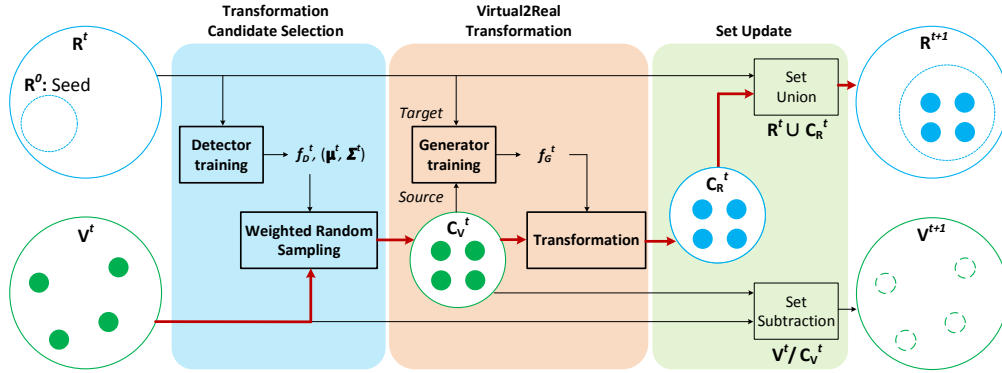


Figure 2: **Progressive Transformation Learning (PTL) pipeline.** The red arrow indicates the processing flow of the virtual images selected to be added to the training set.

on measuring the domain gap, the Mahalanobis distance for  $\mathbf{x}_{new}$  is calculated at multiple image scales, and the minimum distance is used as the domain gap, as follows:

$$d(\mathbf{x}_{new}) = \min_{s \in S} (\{d(\mathbf{x}_{new}^s)\}), \quad (4)$$

where  $\mathbf{x}_{new}^s$  is the resized image of  $\mathbf{x}_{new}$  to be  $s \times s$ .  $S$  is the set of resizing factors. In our experiments, we use  $S = \{128, 256, 384, 512\}$ .

### 3.2 PROGRESSIVE TRANSFORMATION LEARNING

Our objective is to expand the training set consisting of real images by adding virtual images which are transformed to intimate real images. The virtual2real transformation can be performed by a generator trained by treating virtual images and real images as “source” and “target”, respectively, in the conditional GAN framework. Inevitably, the transformation quality of the trained generator is degraded when the domain gap between the source domain and the target domain is large. To prevent the degraded transformation quality due to the large domain gap, we introduce *Progressive Transformation Learning (PTL)*, which progressively and iteratively expands the training set with a subset of virtual images carefully selected to avoid the large domain gap. PTL goes through three steps for each iteration (Fig 2): i) sampling a subset of virtual images from a virtual image pool by giving heavier weights to images close to the current training set (Transformation candidate selection), ii) transforming the selected images to be realistic (Virtual2real transformation), and iii) adding the transformed images to the training set while excluding the selected images from the virtual image pool (Set update). The details of each step are described next.

**Transformation candidate selection.** When selecting transformation candidates, we must consider two conflicting claims simultaneously: i) to suit the purpose of PTL, virtual images with a small domain gap should be selected, but ii) to suit the purpose of expanding the training set, virtual images with appearances that rarely appear in the training set, which usually implies a large domain gap, should also be selected.

To jointly consider these two claims, we use weighted random sampling. The sampling weight takes the exponential term which gives higher weights to virtual images with smaller domain gaps, while introducing one hyper-parameter  $\tau$  to control the amplitude of the weights, as follows:

$$w(\mathbf{x}) = \exp\left(-\frac{d(\mathbf{x})}{\tau}\right), \quad (5)$$

where  $d(\mathbf{x})$  is the Mahalanobis distance, which is used to measure the domain gap of  $\mathbf{x}$  from the current training set (eq 4). Intuitively, using a small  $\tau$  allows a more frequent selection of images with smaller domain gaps by giving them larger weights than using a large  $\tau$ . (We use  $\tau=5.0$  throughout all experiments; the detection accuracy is found to be not very sensitive to  $\tau$ .)

In practice, transformation candidates are selected from virtual image pool through the following four steps: i) training the human detector  $f_D^t$  on the current training set of real images  $R^t$ , ii) calculating  $\mu^t$  and  $\Sigma^t$  on  $R^t$  as in eq. 2, iii) calculating weights for each image in the current set of virtual

images  $V^t$  as in eq. 5, and iv) applying weighted random sampling to  $V^t$  to select a pre-defined number  $n$  of transformation candidates. (We use  $n=100$  throughout all experiments.)

**Virtual2Real transformation.** In line with the goal of this paper to obtain a human detector that can identify humans with diverse appearances captured by a UAV, we design the virtual2real transformation to focus on the person region rather than the background of the selected virtual images. To do so, we crop the person region in the virtual image, apply the transformation only to this region, and segment the transformed person back to the original background. For accurate segmentation, the pixel-wise segmentation mask is required. Obtaining such pixel-wise segmentation mask at no cost is another benefit of using virtual images.

The conditional GAN framework (Mirza & Osindero, 2014), in which the generator is trained to transform a given input image from source styles into target styles, is widely used to transform virtual images to look like real images (Isola et al., 2017; Li & Lee, 2021). Among many variants of conditional GANs, we use CycleGAN (Zhu et al., 2017) in which the generator is trained to minimize the reconstruction error between the input image and the reconstructed image transformed back to the original style of the input image after the initial transformation to the target style. It is shown in Mirza & Osindero (2014) that the transformation with CycleGAN is likely to maintain the original object pose while changing detailed styles such as patterns (e.g., transforming a white horse into a zebra in the same pose). We intend to borrow this characteristic of CycleGAN to transform virtual images in the direction that makes the detailed styles realistic while maintaining the overall human appearances, which depend on various viewing angles or human poses.

In practice, the virtual2real transformation generator  $f_G^t$  is trained using the CycleGAN framework by treating the selected transformation candidates  $C_V^t$  and the current set of real images  $R^t$  as “source” and “target”, respectively. Then,  $C_V^t$  are transformed to realistic transformed images  $C_R^t$  by applying the virtual2real transformation generator.

**Set update.** After the transformed images  $C_R^t$  are acquired from the selected transformation candidates  $C_V^t$ , PTL updates the current real image set  $R^t$  and the current virtual image set  $V^t$  as follows:

$$R^{t+1} = R^t \cup C_R^t \quad V^{t+1} = V^t / C_V^t \quad (6)$$

When this progressive learning is terminated, the final human detection model can be acquired by training on the final set of real images.

In practice, the first two steps of PTL are applied to the tightly cropped image region of human region, but in the ‘set update’ step, the entire image including the human region and the background is added to the training set for training the human detector. More precisely, when training the virtual2real transformation generator, the tightly cropped image region around each human from the training images are used as the “target”.

## 4 EXPERIMENTS

**Datasets and evaluation metrics.** We perform experiments on three real UAV-based datasets, Vis-Drone (Zhu et al., 2021), Okutama-Action (Barekatian et al., 2017), and ICG (ICG), and one virtual dataset, *Archangel-Synthetic* (Shen et al., 2022), all including human instances. *Archangel-Synthetic* consists of various virtual characters with different poses across a range of altitudes and circle radii with different camera pitch angles (i.e., 16.9K images with eight characters, three poses, ten altitudes, six circle radii, and twelve camera pitch angles). Each image in *Archangel-Synthetic* accompanies metadata about the above imaging conditions, allowing us to analyze how the feature distributions of virtual characters evolve with respect to these imaging conditions when PTL progresses. We use AP@.5 and AP@[.5:.95] as evaluation metrics for all experiments.

**Detector.** For the detector, we use RetinaNet (Lin et al., 2017) that uses the feature pyramid network (FPN) to provide a rich multi-scale feature pyramid and processes features at all scale levels with the same subnetwork responsible for the final classification and bounding box regression. It is important to use the same subnetwork across all scale levels since the domain gap for each virtual image should be measured in a shared representation space regardless of the image size. Most other object detectors using FPN (e.g., SSD (Liu et al., 2016) and v4 or later versions of YOLO (Bochkovskiy et al., 2020; Jocher, 2020; Li et al., 2022b; Wang et al., 2022b)) use different subnetworks at each

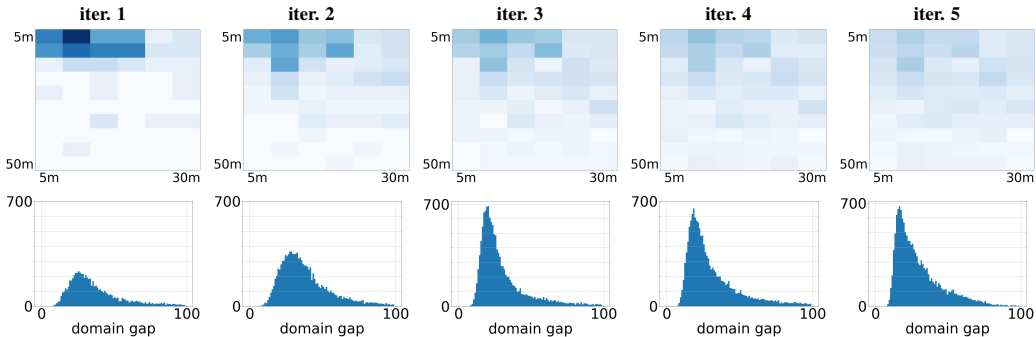


Figure 3: **Analysis of the use of virtual images when PTL progresses.** The figures in the top row show the accumulated distribution of transformation candidates with respect to camera locations (i.e., altitude and rotation circle radius from the target human in  $x$  and  $y$  axes, respectively) for each PTL iteration. Darker bins indicate that more virtual images have been added to the training set. The figures in the bottom row ( $x$  axis: domain gap,  $y$  axis: the corresponding number of virtual images) show the domain gap distribution of virtual images measured by eq. 4. These figures are collected from the experimental setup of using 50 real images from the VisDrone dataset for training.




Figure 4: **Learning curves** of the two metrics (AP@.5 and AP@[.5:.95]) on the three datasets

scale level. However, PTL is not structurally limited to RetinaNet as it can be used with any detector with minor modifications, such as adding one shared layer across all scale levels.

#### 4.1 PROPERTIES OF PROGRESSIVE LEARNING

**Which virtual images are selected for each PTL iteration?** The top row of Figure 3 shows the change in the accumulated distribution of virtual images added to the training set via each PTL iteration with respect to camera locations. By examining the distribution, we can identify which camera locations of the virtual images contribute more to the training set at each PTL iteration.

It is observed that after the 1st PTL iteration, most virtual images included in the training set are taken from the camera locations close to the human subjects. As PTL progresses, the camera locations of the virtual images included in the training set gradually spread across the UAV altitudes and rotation circle radii. Consequently, after the 5th PTL iteration, transformed virtual images with diverse appearances from much broader camera locations are included in the final training set. This demonstrates that the proposed transformation candidate selection process is adequately designed to consider the two conflicting claims together.

**How close does the domain gap get as PTL progresses?** The bottom row of Figure 3 shows the domain gap between the virtual images and the training set at each PTL iteration. We can observe the domain gap distribution of virtual images to the training set gradually becomes narrower and smaller. Additionally, some virtual images which have not been included in the training set also appear in the long tail of the distribution.

**Accuracy variation as PTL evolves.** Figure 4 shows how the accuracy changes as PTL progresses on the three datasets. Overall, for AP@.5, accuracy increases rapidly until the 3rd iteration and does not change significantly thereafter. On the other hand, for AP@[.5:.95], accuracy continues to increase even after the 3rd iteration on the Okutama-Action and ICG datasets. This can be interpreted such that human bounding boxes are estimated more accurately as PTL progresses on these two datasets.

Table 1: **Low-shot learning accuracy** with 20 and 50 real images. AP@.5 and AP@[.5:.95] are reported in each bin. For PTL, the margin from the baseline accuracy is shown below the reported accuracy. The best accuracy for each setting is shown in bold. R and V denote the set of real images and the set of virtual images, respectively.

method	train set	VisDrone				Okutama-Action				ICG			
		20		50		20		50		20		50	
baseline	R	3.74/ 1.09	6.42/ 1.86	41.61/ 11.23	49.84/ 13.76	49.35/ 14.69	66.75/ 23.91						
pretrain-finetune	R+V	4.99/ 1.46	6.25/ 1.99	44.57/ 12.78	49.06/ 15.08	66.92/ 26.67	68.41/ 29.73						
naive merge	R+V	3.41/ 1.02	5.18/ 1.65	34.26/ 9.21	48.33/ 14.61	55.95/ 20.76	65.68/ 26.73						
w/ transform	R+V	1.26/ 0.49	4.02/ 1.37	27.37/ 7.84	41.36/ 12.64	48.02/ 17.62	65.03/ 27.21						
<b>PTL</b> (5th itr.)	R+V	6.83/ 1.94	9.09/ 2.85	52.89/ 15.57	59.90/ <b>18.48</b>	69.11/ <b>27.33</b>	<b>74.14/ 31.41</b>						
		+3.09/+0.85	+2.67/+0.99	+11.28/ +4.34	+10.06/ +4.72	+19.76/+12.64	+7.39/ +7.50						
<b>PTL</b> (best)	R+V	<b>7.52/ 2.13</b>	<b>9.33/ 2.94</b>	<b>53.82/ 15.59</b>	<b>60.65/ 18.48</b>	<b>70.23/ 27.33</b>	<b>74.14/ 31.41</b>						
		+3.78/+1.04	+2.91/+1.08	+12.21/ +4.36	+10.81/ +4.72	+20.88/+12.64	+7.39/ +7.50						

## 4.2 RESULTS ON LOW-SHOT LEARNING

**Baselines.** In this section, we compare PTL with the method utilizing only real images for training (i.e., ‘baseline’) and other three methods also leveraging virtual images in conjunction with real images for training (i.e., ‘pretrain-finetune’, ‘naive merge’, and ‘naive merge w/ transform’) in terms of human detection accuracy. ‘Pretrain-finetune’ is the strategy to train a model by pre-training on virtual images and then fine-tuning on real images, which is the most widely used approach leveraging virtual images in previous works (Handa et al., 2016; Kim et al., 2016; Gaidon et al., 2016; Fabbri et al., 2018; Varol et al., 2018; Mu et al., 2020; Mishra et al., 2022; Guo et al., 2022; Jin et al., 2022; Baek & Shim, 2022). ‘Naive merge’ is the strategy that uses a training set naively merging from real and virtual images for model training, which has also been used in previous works (Ros et al., 2016; Richter et al., 2016). ‘Naive merge w/ transform’ naively adds transformed virtual images to the training set, where the transformation generator is trained with the CycleGAN framework by considering all virtual images and all real images as “source” and “target”.

**Main Results.** In Table 1, we compare PTL to the baselines in terms of human detection accuracy in two low-shot detection regimes (i.e., 20 and 50 real images are used for training) on the three real-world UAV-based datasets. Low-shot detection is a suitable task to validate the proper use of virtual images as notable effects can be expected from adequately expanded datasets.

In all cases, the previous methods leveraging virtual images do not present significantly better, or even worse, accuracy than their counterpart (i.e., ‘baseline’) using real images only for training. ‘Pretrain-finetune’ is the only method, except for PTL, that presents better accuracy than the baseline in some cases. This two-step method is effective in avoiding adverse effects due to the large domain gap while indirectly taking advantage of the large-scale dataset. However, the increase in accuracy is marginal as the task-specific properties (i.e., human detection) of the large-scale virtual dataset used for pre-training are not fully exploited due to the catastrophic forgetting issue inherent in this indirect method. In addition, ‘naive merge w/ transformation’, which is the only previous method that uses transformed virtual images, presents significantly reduced accuracy, which can be regarded as the adverse effects (e.g., transformation quality degradation) when the large domain gap is not properly addressed. The final model obtained after the 5th PTL iteration or the best model achieved when PTL progresses consistently presents significantly better performance than any compared method. This demonstrates *the effectiveness of PTL such that accuracy is substantially improved by expanding the training set using virtual images while the large domain gap is appropriately addressed.*

**Ablation: the effect of  $\tau$ .** In Table 2, we compare three different  $\tau$  values to investigate the effect of  $\tau$  used to control the sampling weights when selecting transformation candidates (eq. 5). In general, the change in accuracy with respect to  $\tau$  is not obvious. Interestingly,  $\tau=10.0$  gives the best accuracy for AP@.5 while  $\tau=2.0$  is the best for AP@[.5:.95]. That is, using a smaller  $\tau$  to select more virtual images with small domain gaps leads to more accurate estimation of the bounding

Table 2: **Varying  $\tau$**  in PTL (best) on the VisDrone dataset with 50-shot learning setup.

$\tau$	2.0	5.0	10.0
AP@.5	9.45	9.33	<b>9.68</b>
AP@[.5:.95]	<b>2.98</b>	2.94	2.89



Table 3: **Cross-domain detection accuracy.** The table shows experiments with  $3 \times 3$  cross-domain setups. For each setup, datasets shown to the left and right of the arrow are the training and test sets, respectively. The accuracies of PTL and the baseline without using virtual images for training are shown. Setups on the top use training and test images from the same domain, which provides a baseline accuracy in the cross-domain setups. All setups in each column are tested on the same dataset and the same low-shot regime.

method	VisDrone				Okutama-Action				ICG			
	20		50		20		50		20		50	
	<i>Vis</i> $\rightarrow$ <i>Vis</i>				<i>Oku</i> $\rightarrow$ <i>Oku</i>				<i>ICG</i> $\rightarrow$ <i>ICG</i>			
baseline	3.74/	1.09	6.42/	1.86	41.61/	11.23	49.84/	13.76	49.35/	14.69	66.75/	23.91
PTL (5th itr.)	6.83/	1.94	9.09/	2.85	52.89/	15.57	59.90/	18.48	69.11/	27.33	74.14/	31.41
PTL (best)	7.52/	2.13	9.33/	2.94	53.82/	15.59	60.65/	18.48	70.23/	27.33	74.14/	31.41
	<i>Oku</i> $\rightarrow$ <i>Vis</i>				<i>Vis</i> $\rightarrow$ <i>Oku</i>				<i>Vis</i> $\rightarrow$ <i>ICG</i>			
baseline	1.62/	0.47	2.04/	0.57	17.13/	4.53	36.82/	9.87	2.92/	0.56	7.46/	1.83
PTL (5th itr.)	2.72/	0.94	3.05/	1.07	30.72/	7.45	42.39/	11.41	26.86/	7.22	29.26/	7.27
PTL (best)	3.00/	1.22	3.56/	1.17	33.25/	8.59	42.39/	11.46	29.60/	7.69	30.01/	7.36
	<i>ICG</i> $\rightarrow$ <i>Vis</i>				<i>ICG</i> $\rightarrow$ <i>Oku</i>				<i>Oku</i> $\rightarrow$ <i>ICG</i>			
baseline	0.54/	0.13	0.99/	0.26	3.56/	0.75	10.27/	2.49	5.37/	1.25	5.23/	1.20
PTL (5th itr.)	1.09/	0.33	1.61/	0.50	11.19/	2.58	14.20/	3.56	28.98/	8.14	25.39/	6.53
PTL (best)	1.58/	1.02	1.70/	0.63	12.82/	2.96	14.20/	3.71	28.98/	8.14	26.62/	6.53

boxes, whereas selecting virtual images with more diverse appearances using a larger  $\tau$  increases the detection accuracy. Taking these two effects into account, we carry out all experiments with  $\tau=5.0$ .

### 4.3 RESULTS ON CROSS-DOMAIN DETECTION

In Table 3, we show the accuracy for the cross-domain setups on the three datasets. The ‘‘cross-domain detection’’ setups are used to validate the impact of using virtual images for training when training and test images have distinct characteristics, such as when the distributions of human appearances with respect to human poses and viewing angles are different.

First, it is observed that using PTL yields much better accuracies than the baseline when using the same real training dataset. Despite the inherent difficulty of cross-domain learning, it is also observed that leveraging virtual images through PTL produces results not far behind the accuracy of the baseline (shown in the first row of Table 3) which uses the training images from the same domain as the test images. However, the improvement in the cross-domain detection accuracy when using the ICG dataset for training is relatively low compared to other cases. This is because the ICG dataset has very different characteristics from other datasets, and virtual images added to the training set do not reduce this difference. Note that if the real images in the initial training set have very distinct appearances, a very different set of virtual images may be selected in the first PTL iteration. Thus, the disparity may not be overcome even if the PTL progresses further. Nevertheless, in general, we can confirm that *human detectors trained using virtual images through PTL can improve substantially detection accuracy, regardless of which real dataset is used during training.*

## 5 DISCUSSION

Our method has been proved effective in leveraging virtual images during training as it presents much better accuracy than any other previous methods for low-shot learning task, where scaling up the training dataset can significantly impact. In addition, compared to the baseline which does not use virtual images, our method also presents remarkable accuracy on cross-domain detection, where the real training and test datasets are from two distinct domains with very different characteristics.

Despite the merit of the proposed method, there is still room to improve PTL further. Specifically, the current version of PTL can only leverage a subset of the entire virtual images within a limited number of iterations, beyond which the accuracy may decrease. We hope that more advanced methods can be developed to address this issue so that PTL can progress for more iterations with a continuous accuracy increase until all virtual images are leveraged for training.

## REFERENCES

- Aerial semantic segmentation drone dataset. URL <http://dronedataset.icg.tugraz.at>.
- Kyungjune Baek and Hyunjung Shim. Commonality in natural images rescues GANs: Pretraining GANs with generic and privacy-free synthetic data. In *Proc. CVPR*, 2022.
- Mohammadamin Barekatin, Miquel Martí, Hsueh-Fu Shih, Samuel Murray, Kotaro Nakayama, Yutaka Matsuo, and Helmut Prendinger. Okutama-action: An aerial view video dataset for concurrent human action detection. In *Proc. CVPRW*, 2017.
- Yoshua Bengio, Pascal Lamblin, Dan Popovici, and Hugo Larochelle. Greedy layer-wise training of deep networks. In *Proc. NeurIPS*, 2006.
- Yoshua Bengio, Jérôme Louradour, Ronan Collobert, and Jason Weston. Curriculum learning. In *Proc. ICML*, 2009.
- Alexey Bochkovskiy, Chien-Yao Wang, and Hong-Yuan Mark Liao. YOLOv4: Optimal speed and accuracy of object detection. arXiv:2004.10934, 2020. URL <https://arxiv.org/abs/2004.10934>.
- Tianyue Cao, Yongxin Wang, Yifan Xing, Tianjun Xiao, Tong He, Zheng Zhang, Hao Zhou, and Joseph Tighe. PSS: Progressive sample selection for open-world visual representation learning. In *Proc. ECCV*, 2022.
- Tianqi Chen, Ian Goodfellow, and Jonathon Shlens. Net2Net: Accelerating learning via knowledge transfer. In *Proc. ICLR*, 2016.
- Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey Hinton. A simple framework for contrastive learning of visual representations. In *Proc. ICML*, 2020.
- Matteo Fabbri, Fabio Lanzi, Simone Calderara, Andrea Palazzi, Roberto Vezzani, and Rita Cucchiara. Learning to detect and track visible and occluded body joints in a virtual world. In *Proc. ECCV*, 2018.
- Scott E. Fahlman and Christian Lebiere. The cascade-correlation learning architecture. In *Proc. NeurIPS*, 1989.
- Adrien Gaidon, Qiao Wang, Yohann Cabon, and Eleonora Vig. Virtual worlds as proxy for multi-object tracking analysis. In *Proc. CVPR*, 2016.
- Linyuan Gong, Di He, Zhuohan Li, Tao Qin, Liwei Wang, and Tie-Yan Liu. Efficient training of BERT by progressively stacking. In *Proc. ICML*, 2019.
- Alex Graves, Marc G Bellemare, Jacob Menick, Remi Munos, and Koray Kavukcuoglu. Automated curriculum learning for neural networks. In *Proc. ICML*, 2017.
- Jean-Bastien Grill, Florian Strub, Florent Altché, Corentin Tallec, Pierre H. Richemond, Elena Buchatskaya, Carl Doersch, Bernardo Avila Pires, Zhaohan Daniel Guo, Mohammad Gheshlaghi Azar, Bilal Piot, Koray Kavukcuoglu, Rémi Munos, and Michal Valko. Bootstrap your own latent - a new approach to self-supervised learning. In *Proc. NeurIPS*, 2020.
- Xi Guo, Wei Wu, Dongliang Wang, Jing Su, Haisheng Su, Weihao Gan, Jian Huang, and Qin Yang. Learning video representations of human motion from synthetic data. In *Proc. CVPR*, 2022.
- Ankur Handa, Viorica Pătrăucean, Vijay Badrinarayanan, Simon Stent, and Roberto Cipolla. SceneNet: Understanding real world indoor scenes with synthetic data. In *Proc. CVPR*, 2016.
- Kaiming He, Haoqi Fan, Yuxin Wu, Saining Xie, and Ross Girshick. Momentum contrast for unsupervised visual representation learning. In *Proc. CVPR*, 2020.
- Judy Hoffman, Eric Tzeng, Taesung Park, Jun-Yan Zhu, Phillip Isola, Kate Saenko, Alexei Efros, and Trevor Darrell. CyCADA: Cycle-consistent adversarial domain adaptation. In *Proc. ICML*, 2018.

- Phillip Isola, Jun-Yan Zhu, Tinghui Zhou, and Alexei A. Efros. Image-to-image translation with conditional adversarial networks. In *Proc. CVPR*, 2017.
- Zhao Jin, Yinjie Lei, Naveed Akhtar, Haifeng Li, and Munawar Hayat. Deformation and correspondence aware unsupervised synthetic-to-real scene flow estimation for point clouds. In *Proc. CVPR*, 2022.
- Glenn Jocher. ultralytics/yolov5: v3.1 - Bug Fixes and Performance Improvements. <https://github.com/ultralytics/yolov5>, 2020. URL <https://doi.org/10.5281/zenodo.4154370>.
- Tero Karras, Timo Aila, Samuli Laine, and Jaakko Lehtinen. Progressive growing of GANs for improved quality, stability, and variation. In *Proc. ICLR*, 2018.
- Prannay Kaul, Weidi Xie, and Andrew Zisserman. Label, verify, correct: A simple few shot object detection method. In *Proc. CVPR*, 2022.
- Faisal Khan, Bilge Mutlu, and Jerry Zhu. How do humans teach: On curriculum learning and teaching dimension. In *Proc. NeurIPS*, 2011.
- Dahun Kim, Sanghyun Woo, Joon-Young Lee, and In So Kweon. Video panoptic segmentation. In *Proc. CVPR*, 2016.
- M. Pawan Kumar, Benjamin Packer, and Daphne Koller. Self-paced learning for latent variable models. In *Proc. NeurIPS*, 2010.
- Kimin Lee, Kibok Lee, Honglak Lee, and Jinwoo Shin. A simple unified framework for detecting out-of-distribution samples and adversarial attacks. In *Proc. NeurIPS*, 2018.
- Changlin Li, Bohan Zhuang, Guangrun Wang, Xiaodan Liang, Xiaojun Chang, and Yi Yang. Automated progressive learning for efficient training of vision transformers. In *Proc. CVPR*, 2022a.
- Chen Li and Gim Hee Lee. From synthetic to real: Unsupervised domain adaptation for animal pose estimation. In *Proc. CVPR*, 2021.
- Chuyi Li, Lulu Li, Hongliang Jiang, Kaiheng Weng, Yifei Geng, Liang Li, Zaidan Ke, Qingyuan Li, Meng Cheng, Weiqiang Nie, Yiduo Li, Bo Zhang, Yufei Liang, Linyuan Zhou, Xiaoming Xu, Xiangxiang Chu, Xiaoming Wei, and Xiaolin Wei. YOLOv6: A single-stage object detection framework for industrial applications. arXiv:2209.02976, 2022b. URL <https://arxiv.org/abs/2209.02976>.
- Tsung-Yi Lin, Priya Goyal, Ross Girshick, Kaiming He, and Piotr Dollár. Focal loss for dense object detection. In *Proc. ICCV*, 2017.
- Qing Liu, Adam Kortylewski, Zhishuai Zhang, Zizhang Li, Mengqi Guo, Qihao Liu, Xiaoding Yuan, Jiteng Mu, Weichao Qiu, and Alan Yuille. Learning part segmentation through unsupervised domain adaptation from synthetic vehicles. In *Proc. CVPR*, 2022.
- Wei Liu, Dragomir Anguelov, Dumitru Erhan, Christian Szegedy, Scott Reed, Cheng-Yang Fu, and Alexander C. Berg. SSD: Single shot multibox detector. In *Proc. ECCV*, 2016.
- Prasanta Chandra Mahalanobis. On the generalised distance in statistics. *Proc. National Institute of Sciences of India.*, 1936.
- Mehdi Mirza and Simon Osindero. Conditional generative adversarial nets. arXiv:1411.1784, 2014. URL <https://arxiv.org/abs/1411.1784>.
- Samarth Mishra, Rameswar Panda, Cheng Perng Phoo, Chun-Fu (Richard) Chen, Leonid Karlinsky, Kate Saenko, Venkatesh Saligrama, and Rogerio S. Feris. Task2Sim: Towards effective pre-training and transfer from synthetic data. In *Proc. CVPR*, 2022.
- Jiteng Mu, Weichao Qiu, Gregory Hager, and Alan Yuille. Learning from synthetic animals. In *Proc. CVPR*, 2020.

- Haibo Qiu, Baosheng Yu, Dihong Gong, Zhifeng Li, Wei Liu, and Dacheng Tao. SynFace: Face recognition with synthetic data. In *Proc. ICCV*, 2021.
- Stephan R. Richter, Vibhav Vineet, Stefan Roth, and Vladlen Koltun. Playing for data: Ground truth from computer games. In *Proc. ECCV*, 2016.
- German Ros, Laura Sellart, Joanna Materzynska, David Vazquez, and Antonio M. Lopez. The SYNTHIA dataset: A large collection of synthetic images for semantic segmentation of urban scenes. In *Proc. CVPR*, 2016.
- Yi-Ting Shen, Yaesop Lee, Heesung Kwon, Damon M. Conover, Shuvra S. Bhattacharyya, Nikolas Vale, Joshua D. Gray, G. Jeremy Leong, Kenneth Evensen, and Frank Skirlo. Archangel: A hybrid UAV-based human detection benchmark with position and pose metadata. arXiv:2209.00128, 2022. URL <https://arxiv.org/abs/2209.00128>.
- Leslie N. Smith, Emily M. Hand, and Timothy Doster. Gradual DropIn of layers to train very deep neural networks. In *Proc. CVPR*, 2016.
- Valentin I. Spitzkovsky, Hiyan Alshawi, and Daniel Jurafsky. Baby steps: How “less is more” in unsupervised dependency parsing. In *Proc. NeurIPS Workshop*, 2009.
- Gul Varol, Javier Romero, Xavier Martin, Naureen Mahmood, Michael J. Black, Ivan Laptev, and Cordelia Schmid. Learning from synthetic humans. In *Proc. CVPR*, 2018.
- Can Wang, Sheng Jin, Yingda Guan, Wentao Liu, Chen Qian, Ping Luo, and Wanli Ouyang. Pseudo-labeled auto-curriculum learning for semi-supervised keypoint localization. In *Proc. ICLR*, 2022a.
- Chien-Yao Wang, Alexey Bochkovskiy, and Hong-Yuan Mark Liao. YOLOv7: Trainable bag-of-freebies sets new state-of-the-art for real-time object detectors. arXiv:2207.02696, 2022b. URL <https://arxiv.org/abs/2207.02696>.
- Guangcong Wang, Xiaohua Xie, Jianhuang Lai, and Jiakuan Zhuo. Deep growing learning. In *Proc. ICCV*, 2017.
- Tao Wei, Changhu Wang, Yong Rui YONGRUI, and Chang Wen Chen. Network morphism. In *Proc. ICML*, 2016.
- Qi Yan, Jianhao Zheng, Simon Reding, Shanci Li, and Iordan Doytchinov. CrossLoc: Scalable aerial localization assisted by multimodal synthetic data. In *Proc. CVPR*, 2022.
- Minjia Zhang and Yuxiong He. Accelerating training of transformer-based language models with progressive layer dropping. In *Proc. neurIPS*, 2020.
- Jun-Yan Zhu, Taesung Park, Phillip Isola, and Alexei A. Efros. Unpaired image-to-image translation using cycle-consistent adversarial networks. In *Proc. ICCV*, 2017.
- Pengfei Zhu, Longyin Wen, Dawei Du, Xiao Bian, Heng Fan, Qinghua Hu, and Haibin Ling. Detection and tracking meet drones challenge. *IEEE Trans. Pattern Anal. Mach. Intell.*, 2021.