

# One for More: Selecting Generalizable Samples for Generalizable ReID Model

Enwei Zhang<sup>1</sup>, Xinyang Jiang<sup>1†\*</sup>, Hao Cheng<sup>1</sup>, Ancong Wu<sup>2</sup>, Fufu Yu<sup>1</sup>, Ke Li<sup>1</sup>,  
Xiaowei Guo<sup>1</sup>, Feng Zheng<sup>3</sup>, Weishi Zheng<sup>2</sup>, Xing Sun<sup>1†</sup>

<sup>1</sup> Tencent Youtu Lab, Shanghai, China

<sup>2</sup> Sun Yat-sen University, Shenzhen, China

<sup>3</sup> Southern University of Science and Technology, Shenzhen, China

## Abstract

Current training objectives of existing person Re-Identification (ReID) models only ensure that the loss of the model decreases on selected training batch, with no regards to the performance on samples outside the batch. It will inevitably cause the model to over-fit the data in the dominant position (e.g., head data in imbalanced class, easy samples or noisy samples). The latest resampling methods address the issue by designing specific criterion to select specific samples that trains the model generalize more on certain type of data (e.g., hard samples, tail data), which is not adaptive to the inconsistent real world ReID data distributions. Therefore, instead of simply presuming on what samples are generalizable, this paper proposes a one-for-more training objective that directly takes the generalization ability of selected samples as a loss function and learn a sampler to automatically select generalizable samples. More importantly, our proposed one-for-more based sampler can be seamlessly integrated into the ReID training framework which is able to simultaneously train ReID models and the sampler in an end-to-end fashion. The experimental results show that our method can effectively improve the ReID model training and boost the performance of ReID models.

## Introduction

Person re-identification (ReID) aims at recognizing pedestrians across non-overlapping camera views, which increasingly draws attention due to its wide applications in surveillance, tracking, smart retail, etc (Zheng et al. 2015; Zhao et al. 2017; Sun et al. 2018). As deep learning prevails, the CNN based ReID methods progress rapidly and achieve impressive performance on benchmark datasets. However, ReID remains a challenging problem due to the view variance, domain changes, partial occlusion and many other factors. Most of the existing ReID methods solve these issues by designing new network structures or innovative loss functions (Sun et al. 2018; Wang et al. 2018b), but very few pay attention to the nature of training data distribution.

\*Enwei Zhang and Xinyang Jiang contribute equally

†Corresponding Author (winfredsun@tencent.com, xinyangj@zju.edu.cn)

Copyright © 2021, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

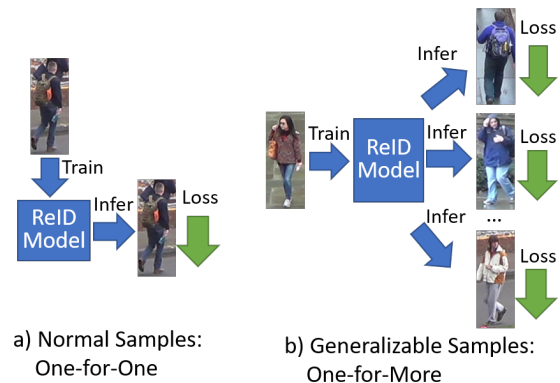


Figure 1: Comparison between one-for-one objective and one-for-more objective.

In standard ReID training task, at each iteration, after updating the model with a batch of selected samples, a normal training objective only makes sure the loss of these sample decreases, but has no guarantee that the loss of the samples outside the batch will also decreases, as shown in Figure 1 a). We call this type of objectives one-for-one objectives because the samples are selected for improving model performance on themselves, with no regards to performance on the samples outside the batch. As shown in 1 b), after updating model with the selected sample, if the loss of more samples decreases other than the sample itself, we call this sample a more generalizable sample. Intuitively, training with generalizable samples gives a model with more generalization ability.

Real world ReID data distribution is usually highly imbalanced with notable amount of hard samples (e.g., view/pose change, occlusion, camera resolution etc) and noisy data, and one-for-one objective could potentially cause models over-fitting on a dominate part of the data. To avoid over-fitting, methods including re-sampling and re-weighting propose different criteria to select samples to train the model to generalize more on certain type of data. For example, following are some typical selection criteria for state-of-the-art ReID methods:

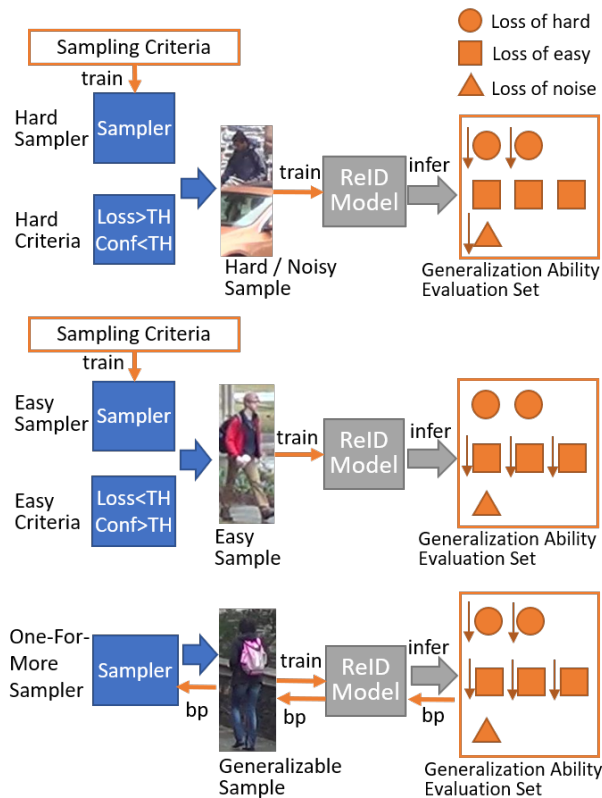


Figure 2: Comparisons of our methods with different sampling methods.

- **Select Easy Sample** Methods like curriculum learning (Wang et al. 2018a)(Guo et al. 2018) pre-define noisy labels as samples with low training loss or high prediction confidence and proposes to emphasize on these samples to prevent model over-fitting on noisy labels.
- **Select Hard Sample** Methods like hard sample mining (Hermans\*, Beyer\*, and Leibe 2017) and focal loss (Lin et al. 2017) define hard samples as samples with large training loss or low prediction confidence, and proposes to emphasize on these samples to help model generalize better on samples difficult to recognize.
- **Learn Easy Sampler** Methods like self-paced Learning (Zhou et al. 2018) and abstention (Thulasidasan et al. 2019) do not directly use a selection criteria to select samples. Instead, the criteria becomes an objective function to train an automatic data sampler. For example, Abstention trains an easy data sampler to select samples with the minimum training loss.
- **Learn Hard Sampler** Similarly to easy data sampler, methods like DE-DSP(Duan et al. 2019) train an automatic hard data sampler by maximizing the the training loss of selected samples.

Besides above-mentioned methods, there are also other types of selection criteria making model generalize on different types of samples, such as tail data (up-sampling or

over-sampling), image with certain attributes (Wu et al. 2018).

All the above-mentioned methods are based on selection criterion that selects specific samples to train model to generalize on certain type of data. Thus, these methods are usually not adaptable when training on another type of data distribution. For example, as shown in Figure 2, hard sampling causes over-fitting on noisy labeled data which also has large training loss, and easy sampling causes over-fitting on easy samples that takes dominate place in training set. As a result, instead of giving unreliable presumption of what is good criterion for selecting generalizable sample, we propose a new objective and criterion that directly takes sample’s ability to make model generalize on more data as a loss function, called one-for-more objective. As shown in last row of Figure 2. A batch of samples are first selected to update the ReID model. Then, another evaluation set of samples outside the training batch is sampled, and the generalization ability of the selected training batch is evaluated by computing the updated model’s loss on this set. Thus, the data sampler is trained by minimizing the training loss outside the selected training batch, which is to back-propagate the training loss. The detail implementation will be introduced in section 3.

The one-for-more sampler training process is integrated into the ReID model training process, forming an end-to-end framework, which at each iteration, simultaneously updates the data sampler and updates the ReID model with samples selected by the sampler. To make the data sampler more suitable for ReID, besides sampling individual images, we propose a pair-wise data sampler that considers correlation between images and support pair-wise sample selection for contrastive loss.

In conclusion, to avoid over-fitting caused by standard one-for-one objective in a complex real world data distribution, this paper proposes an one-for-more objective based end-to-end framework that simultaneously trains ReID model with selected samples and the data sampler. Instead of optimizing sample itself, our method is able to select generalizable samples that update the ReID model towards generalizing on samples outside the training batch.

## Related Work

As mentioned in the last section, to prevent over-fitting on small sub-set of the training data caused by one-for-one objectives, over the years, many sampling based works have been proposed to adapt different types of data distribution.

**Imbalance Identities.** Re-sampling/re-weighting (He and Garcia 2009; Chawla et al. 2002; Shrivastava, Gupta, and Girshick 2016) based methods are one of the most widely used strategy to tackle imbalance problem including up-sampling the minority samples or down-sampling the majority samples. Other works mainly focus on adjusting the model or loss function. For example, (Zhang et al. 2017) solves the imbalance problem between relevant pairs and irrelevant pairs of samples by introducing a range loss that defines the sample relevance at identity level. (Zhong et al. 2019) propose to treats head data and tail data in different way , and apply two different sets of loss functions. (Liu et al. 2019) propose an external memory structure to help

learn better features on tail samples, but their method is for classification task, which is not compatible with recognition and retrieval.

**Hard Sample Mining.** One of the goals of re-sampling and re-weighting methods is to locate and emphasize the learning on samples hard to converge (e.g., samples with small inter-class variation and large intra-class variation). Adaboost (Freund and Schapire 1997) is a widely used machine learning algorithm, which iteratively trains new models on hard samples found by old models. Online hard negative mining (Shrivastava, Gupta, and Girshick 2016) and focal loss (Lin et al. 2017) propose loss based hard sample mining method that selects hard training samples for object detection. (Hermans\*, Beyer\*, and Leibe 2017) proposes a easy-to-implement online hard mining method for triplet loss and achieves great performance improvement in person re-identification. (Schroff, Kalenichenko, and Philbin 2015) propose to only select negative samples that further away from the anchor than the positive exemplar, but the distance to anchor still close to the anchor-positive distance (i.e. semi-hard samples). DE-DSP(Duan et al. 2019) and Meta-weight Net (Shu et al. 2019; Ren et al. 2018) propose to learn re-weighting functions that automatically assign importance weight to training samples, which is somewhat similar to our method. Compared to these works, our methods propose a one-for-more objective which able to select samples that update model toward generalize on the entire training set.

**Noise Resist.** Most of the person re-identification datasets contain a small portion of noisy labels. Many studies propose to strengthen the resistance toward noisy labels by customized sampling strategy (Kumar, Packer, and Koller 2010), network structure (Yu et al. 2019), data cleaning (Tanaka et al. 2018) or data augmentation (Srivastava et al. 2014; Zhang et al. 2018). The experiments will show that our method is robust to noisy labels as long as the noisy level is below 20% (consistent with most of the ReID data distribution).

## The Proposed Method

Our method trains two neural network models, namely the ReID model and the data sampler. The data sampler is responsible for outputting a sampling policy to select samples from the training set and then the ReID model is updated with the selected training batch. As shown in 3, there are two learning branches in our proposed method, which are jointly trained in an end-to-end fashion.

The ReID model learning process is a typical deep learning schema. At each iteration, a mini-batch is drawn from training set with the sampling policy (i.e. sampling probability of each sample) generated by data sampler and fed into the ReID model to obtain the training loss. The gradient of the training loss is then back propagated to update the model weights.

The data sampler learning process involves two model update processes in one iteration. First, a sample batch selected by the sampler is used to update the ReID model. Then, another sampler training batch is drawn from the original training set distribution to compute the one-for-more sampler training loss with the updated ReID model. Finally, one-

for-more loss is back propagated all the way through the ReID update operation to the data sampler. Since data sampling process is non-derivable, instead of directly sampling a batch, we approximate an expectation of the training batch sampled with current data sampler.

We elaborate on the detailed implementation of the two learning branches in the following sections.

## Background

For a person re-identification problem, given any image denoted as  $x$ , we need to train a ReID model  $h$  to predict its corresponding identity  $y$ .

Deep learning based ReID model uses mini-batch stochastic gradient descent algorithm to solve the ReID loss function. At each iteration, a small sub-set of samples are drawn from the training set for model update. Denoting the sampling policy as  $\hat{P}$ , at each iteration, the model weights is updated with the mini-batch drawn with the sampling policy:

$$w^{(t+1)} = w^{(t)} - \beta \nabla_{w^{(t)}} \frac{1}{K} \sum_{x_i \sim \hat{P}} L(h(x_i, w^{(t)}), y_i), \quad (1)$$

where,  $w^{(t)}$  is the parameters of the ReID model at  $t$ -th iteration,  $K$  is the size of the mini-batch, and  $L$  is the ReID Loss function. The goal of our method and other re-sampling methods is to find an optimal sampling policy that selects generalizable samples for ReID model update.

## One-for-more Objective for Data Sampler

**Data Sampler Modeling** For most of the existing re-sampling methods, sampling policy  $\hat{P}$  is not automatically learned, but a predefined fixed strategy. For example, online hard sample mining defines a policy that first sampling mini-batch from training data from a uniform probability and then selects triplet samples with lowest similarity in from this batch.

Instead of a predefined strategy, we propose to learn a data sampler that predicts the probability  $\hat{P}(x|h)$  to select sample  $x$  given current ReID model  $h$ . By choosing Multi-layer perception as an energy function, the probability for sampler to select an image  $x$  is modeled as:

$$\hat{P}(x|h) = \frac{e^{MLP(h(x,w),\theta)}}{\sum_{i=1}^N e^{MLP(h(x_i,w),\theta)}}, \quad (2)$$

where  $\theta$  is the set of parameters of probability function  $\hat{P}$  and  $N$  is the size of the dataset.

The above-mentioned sampling policy only considers visual information of single images and is not suitable for sampling image pairs or triplets. We further extend the proposed data sampler to predict sampling policy based on mutual information of sample pairs. We model the probability to select any image  $x$  given an anchor image  $x_a$  and ReID model  $h$  is modeled as follows:

$$\hat{P}(x|x_a, h) = \frac{e^{MLP([h(x,w), h(x_a,w)], \theta)}}{\sum_{i=1}^N e^{MLP([h(x_i,w), h(x_a,w)], \theta)}}, \quad (3)$$

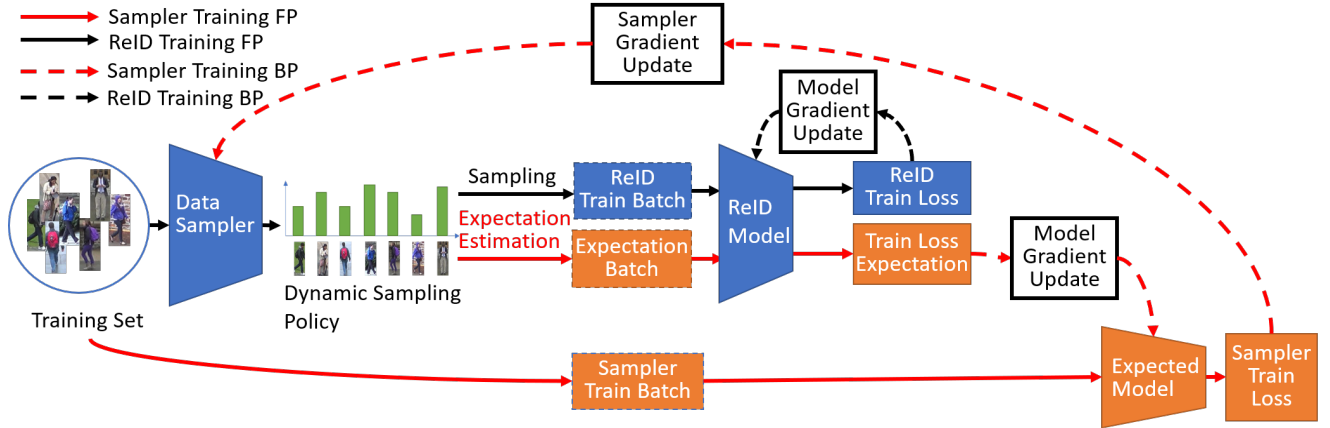


Figure 3: The detailed workflow of ReID method with proposed dynamic sampling policy learning.

where  $[h(x_i, w), h(x_a, w)]$  is the concatenation of feature.

In conclusion, we sample individual image batches to optimize cross entropy based loss with single image sampler  $\hat{P}(x|h)$ . For triplet loss, we first sample anchors with single image sampler, and then given an anchor image, we sample corresponding positive and negative sample with pair-wise sampler  $\hat{P}(x|x_a, h)$  to form a image triplet.

**Data Sampler Optimization** As shown in Figure 3, the optimization of data sampler involves two model update processes. The model is first updated with a batch of samples selected by the trained sampler, and then the objective function for the data sampler is to minimize the ReID loss of the updated model over another set of data.

In detail, since the process of sampling data from a distribution is non-derivable, we use a expectation approximation to replace the actual data sampling. At  $t$ -th iteration, we approximate the expectation of the updated model's weights, given current sampling policy  $\hat{P}$ , denoted as  $E_{\hat{P}}[w^{(t+1)}]$ . Then the optimal policy  $\hat{P}$  is obtained by minimizing the updated model's training loss over all  $N$  training samples:

$$\begin{aligned} \hat{P} &= \arg \min_{\hat{P}} R'(\hat{P}) \\ &= \arg \min_{\hat{P}} \sum_{i=1}^N L(h(x_i, E_{\hat{P}}[w^{(t+1)}]), y_i). \end{aligned} \quad (4)$$

Next we introduce how to obtain the above-mentioned updated model weight expectation  $E_{\hat{P}}[w^{(t+1)}]$ . Assuming at each learning iteration, we draw  $n$  single images  $x_i$  from training set containing  $N$  images. Then the expected number of sample  $x_i$  occurs in the drawn mini-batch is  $n\hat{P}(x_i|h)$ , and then given the generated sampling policy  $\hat{P}$ , the expectation of the ReID model's loss on the drawn batch is:

$$E_{\hat{P}}[L(h(x), y)] = \sum_{i=1}^N \hat{P}(x_i) L(h(x_i, w^{(t)}), y_i), \quad (5)$$

Similarly, the triple loss of data sampled by the image pair

sampler is formulated as:

$$\begin{aligned} E_{\hat{P}}[L(h(x), y)] &= \sum \hat{P}(x_a, x_p, x_n|h) L_{tri} \\ &= \sum \hat{P}(x_a|h) \hat{P}(x_p|x_a, h) \hat{P}(x_n|x_a, h) L_{tri}, \end{aligned} \quad (6)$$

where  $L_{tri}$  is short for:

$$\begin{aligned} L_{tri} &= \max \left( d(h(x_a, w^{(t)}), h(x_p, w^{(t)})) \right. \\ &\quad \left. - d(h(x_a, w^{(t)}), h(x_n, w^{(t)})) + m, 0 \right), \end{aligned} \quad (7)$$

where  $x_a$  is the anchor image;  $x_p$  is the positive sample which has the same identity with anchor image;  $x_n$  is the negative sample of the triplet;  $m$  is the margin value. Following the mini-batch based SGD, it is infeasible to obtain loss expectation based on the entire training set. We draw a subset containing  $K$  samples to approximate the loss expectation, where the values of the samples' selecting probability is normalized to sum to 1.

Due to the page limit, in rest of this section, we take single image sampler  $\hat{P}(x|h)$  as an example to introduce the proposed sampling policy learning methods. Now that we obtain the train loss expectation, we use a common SGD method to update the ReID model. By taking the derivative of the loss expectation, we obtain the expected gradient of the loss with respect to the model weight, noted that the expectation is approximated on a batch containing  $K$  samples:

$$E_{\hat{P}}[\nabla_{w^{(t)}} L(h(x), y)] = \sum_{i=1}^K \hat{P}(x_i|h) \nabla_{w^{(t)}} L(h(x_i, w^{(t)}), y_i). \quad (8)$$

As a result, the expected model weight is updated as:

$$E_{\hat{P}}[w^{(t+1)}] = w^{(t)} - \beta \sum_{i=1}^K \hat{P}(x_i|h) \nabla_{w^{(t)}} L(h(x_i, w^{(t)}), y_i). \quad (9)$$

Finally we evaluate the training loss of the updated model over all samples in the training set. Given a sampler training batch with  $M$  samples is randomly sampled, we substitute Eq.(9) to the optimization objective in Eq.(4) and obtain the

training loss expectation over  $M$  samples, and compute the gradient of this loss function with respect to data sampler’s weight  $\theta$ . Given the gradient, the weights of the data sampler are updated as follows:

$$\theta^{(t+1)} = \theta^{(t)} - \alpha \sum_{i=1}^M \nabla_{\theta} L\left(h(x_i, E_{\hat{P}}(w^{(t+1)})), y_i\right) \quad (10)$$

## ReID Model Training with Automatic Data Sampler

Our training framework uses both cross entropy loss and triplet loss, which is two types of loss function achieves state-of-the-art performance. Our framework learns a single image sampler  $\hat{P}(x|h)$  to sample individual images and learns an image pairs sampler  $\hat{P}(x|x_a, h)$  to sample image pairs. We first use single image sampler to draw anchor images, and then use image pair sampler  $\hat{P}(x|x_a, h)$  to sample positive and negative image pairs given the anchor image.

Directly obtaining the probability of image pairs in  $N^2$  space is infeasible during training. Thus, given the anchor image  $x_a$ , we design a semi-global sampling process to approximate the above-mentioned pair-wise sampling. First we use single image sampler  $\hat{P}(x|h)$  to sample a positive candidate set and a negative candidate set globally from the entire training set, and then the positive image and negative images are sampled from the candidate set with  $\hat{P}(x|x_a, h)$  respectively. We summarize the entire pipeline of our proposed training framework step-by-step in Algorithm 1.

---

### Algorithm 1 Person Re-Identification with Dynamic Data Sampler Learning

---

**Require:** Training data  $D_{tr}$   
**Ensure:** ReID model parameter  $w^{(T)}$   
Initialize  $w^{(0)}$  and  $\theta^{(0)}$ .  
**for**  $t = 0$  **to**  $T - 1$  **do**  
  Sample anchor images  $x_a \sim \hat{P}(x|h)$   
  **for each**  $x_a$  **do**  
    Sample positive candidate set  $X_p$   
    Sample negative candidate set  $X_n$   
    Sample positive pair  $x_p$  from  $X_p$  with  $\hat{P}(x_n|x_a, h)$   
    Sample negative pair  $x_n$  from  $X_n$  with  $\hat{P}(x_n|x_a, h)$   
    Add  $(x_a, x_p, x_n)$  into ReID training batch  
  **end for**  
  Update ReID model weights  $w^{(t+1)}$  with Eq. (1) using ReID training batch  
  Sample expectation evaluation batch  
  Obtain weight expectation of the data sampler with expectation evaluation batch by Eq. (9)  
  Sample data sampler training batch  
  Update weights of data sampler  $\hat{P}$  with sampler training batch by Eq. (10)  
**end for**

---

## Experiments

### Dataset and Evaluation Metrics

Our experiments are conducted on three widely used ReID benchmark datasets.

**Market-1501** dataset contains 32,668 person images of 1,501 identities captured by six cameras. Training set is composed of 12,936 images of 751 identities while testing data is composed of the other images of 750 identities.

**MSMT-17** dataset contains 124,068 person images of 4,101 identities captured by 15 cameras. Training set contains 30,248 images of 1,041 identities while testing set contains 3060 identities.

**DukeMTMC-reID** dataset contains 36,411 person images of 1,404 identities captured by eight cameras, with training set containing 702 identities and 702 identities as the testing set.

**Evaluation Metrics.** Mean average precision (mAP) and matching accuracy (Rank-1) are adopted.

### Implementation Details

The input image size is set to  $256 \times 128$ . The ResNet-50 model with the pretrained parameters on ImageNet is chosen as the backbone network. Common data augmentation include horizontal flipping, random cropping, padding, random erasing (with a probability of 0.5) are used. We adopt Adam optimizer to train our model and set weight decay  $5 \times 10^{-4}$ . The total number of epoch is 200 and the epoch milestones are 50, 100, 160. The learning rate is initialized to  $3.5 \times 10^{-5}$  and is decayed by a factor of 0.1 when the epoch get the milestones. At the beginning, we warm up the models for 10 epochs and the learning rate grows linearly from  $3.5 \times 10^{-5}$  to  $3.5 \times 10^{-4}$ . For data sampler, we choose a fully connect layer as energy function.

### Performance Comparison

To verify the effectiveness of our model, we compare our method with some of the state-of-art re-sampling and re-weighting methods on Market-1501, DukeMTMC and MSMT-17. We further conduct experiments to evaluate the robustness of our method under different types of distribution property. To evaluate robustness on class imbalance, we compare the performance of our methods on dataset with different degree of class imbalance. To evaluate the ability to converge on hard samples (i.e. high intra-identity variation and low inter-identity variation samples), we compare the performance on cross-viewpoint test set.

**Overall Performance Comparison** We compare our method with state-of-art re-sampling and re-weighting method on three benchmark datasets.

The performance comparison of both image sampling methods for cross entropy loss and triplet sampling methods for triplet loss are reported. For a fair comparison, following methods run under the same experiment setting as our proposed methods:

- **Strong Baseline** (Luo et al. 2019). We choose this method as a baseline. To compare the performance of methods sampling single images, the baseline for image sampling methods does not use triplet loss and triplet sampling.

Category	Method	Market-1501		DukeMTMC-reID		MSMT-17	
		mAP	Rank-1	mAP	Rank-1	mAP	Rank-1
Sampling Images	Baseline (CE) (2019)	78.73	91.98	70.12	83.89	46.75	72.25
	Data Prior Distribution <sup>†</sup> (2018)	65.87	86.90	53.42	72.83	-	-
	Focal Loss (2017)	79.03	90.50	67.76	81.46	44.32	68.91
	Self-paced Learning (2018)	80.61	91.78	63.07	77.29	50.01	74.28
	Meta-Weight Net (2019)	79.04	91.54	70.10	82.90	45.43	71.40
	<b>Our Methods</b>	<b>82.27</b>	<b>93.02</b>	<b>73.33</b>	<b>85.86</b>	<b>51.60</b>	<b>75.09</b>
Sampling Triplets	Baseline (OHEM) (2019)	85.7	94.1	75.9	86.2	50.33	74.05
	Triplet loss OHEM <sup>†</sup> (2017)	69.14	84.92	-	-	-	-
	MVP Loss <sup>†</sup> (2019)	80.5	91.4	70.0	83.4	-	-
	Mancs <sup>†</sup> (2018a)	82.3	93.1	71.8	84.9	-	-
	SemiHard OEM (2015)	85.66	93.94	75.83	86.67	49.93	73.44
	<b>Our Methods</b>	<b>87.94</b>	<b>94.89</b>	<b>78.58</b>	<b>88.96</b>	<b>54.69</b>	<b>78.35</b>

Table 1: Performance (%) comparisons to the state-of-the-art results on Market-1501, DukeMTMC-reID and MSMT-17. †: result quote from original paper.

Few-shot Samples	Method	Market-1501		
		mAP	Rank-1	Rank-5
90%	Baseline	71.07	87.29	95.64
	Up-sample	74.63	89.28	96.2
	This work	<b>81.57</b>	<b>92.70</b>	<b>97.36</b>
80%	Baseline	76.95	90.38	96.79
	Up-sample	78.44	91.06	96.97
	This work	<b>83.88</b>	<b>93.79</b>	<b>97.83</b>
70%	Baseline	79.80	91.63	<b>97.12</b>
	Up-sample	80.50	91.75	97.51
	This work	<b>84.68</b>	<b>93.56</b>	<b>97.92</b>

Table 2: Performance (%) comparisons of the class-balance resist method and our method on Market-1501 with different level of class imbalance.

- **Focal loss** (Lin et al. 2017). This is a hard sampling method where the identity loss of each sample is re-weighted with a focal loss weight function.
- **Self-paced Learning** (Zhou et al. 2018). This is a method that learns easy sampler where the identity loss of the strong baseline method is replaced with self-paced learning based loss.
- **Meta-Weight Net** (Shu et al. 2019). This is a method that learn hard sampler where the identity loss of the samples is re-weighted by a weighting function automatically obtained by meta-learning.
- **SemiHard OEM** (Schroff, Kalenichenko, and Philbin 2015). This is a method that select hard samples where the OHEM is replaced with Semi-Hard sampling strategy.

As the source code is not available for some of the state-of-the-art methods (Wu et al. 2018; Wang et al. 2018a; Sun et al. 2019), we quote the results of these methods directly from the original paper, where Mancs (Wang et al. 2018a) is a curriculum learning method to select easy samples, MVP loss (Sun et al. 2019) is a hard sample selecting method and Data Prior Distribution select samples based on attributes (Wu et al. 2018). These methods use similar network structure and loss function to produce relatively fair and compa-

table results to our methods. As shown in Table 1, for both sampling images and triplets, our method outperforms state-of-the-art methods, which verifies the advantage of one-for-more based data sampler for ReID.

We observe that some of the predefined re-weighting/re-sampling methods (e.g. ohem) achieve better performance compared to the uniform sampling baseline, while others do not improve the performance (e.g., semihard ohem). These results show that the predefined re-weighting/re-sampling methods do not generalize on all types of data distribution and extra work is needed to customize selection strategy. Furthermore, some of the methods only achieve performance improvement on one or two of the three datasets and achieve worse performance than baseline on the others. This result further proves that predefined strategy does not generalize well on different datasets.

We observe that our method outperforms the meta-weight net, which also automatically learns a local re-weighting function for cross entropy loss. This results show that learning global sampling policy and pairwise data sampler instead of re-weighting individual image on local batch is more suitable for person ReID task and effectively boosts the performance.

**Performance on Imbalanced Dataset** To verify the effectiveness of our method on imbalanced dataset, we manually enhance the level of identity balance of the existing dataset. For a public ReID dataset, we first sort the identities of the existing dataset based the number of images in them. Then, we down-sample the top- $m$  identities to  $n$  samples each ( $n = 5$  in our experiments). Table 2 shows the performance of baseline methods and our method on imbalance dataset with different level of identity imbalance. We choose a classic up-sampling strategy as the comparison method, which has proven to be a effective method for imbalanced dataset. We observe that, when there is a large amount of tail data, our method significant outperforms the baseline and existing up-sampling strategy.

**Evaluation on Hard Samples** Here we construct a new hard case dataset to verify our method’s ability to con-

Query	Method	mAP	Rank-1	Rank-5
$D_{front}$	baseline	81.49	89.82	95.99
	Our method	<b>84.82</b>	<b>91.50</b>	<b>96.63</b>
$D_{side}$	baseline	82.17	90.56	95.52
	Our method	<b>84.47</b>	<b>91.10</b>	<b>96.33</b>
$D_{back}$	baseline	82.82	91.42	97.03
	Our method	<b>84.82</b>	<b>93.73</b>	<b>97.03</b>

Table 3: Performance Comparison (%) of baseline method and our method on cross-viewpoint retrieval test set.



Figure 4: High probability samples in generalized data distribution at different training stages

verge on datasets with high intra-identity variation and inter-identity variation. Specifically, samples in Market-1501 are classified into three view-point categories: *front*, *side*, *back*. For each query in each test case, to increase the inner-class variation, we only select the images with different views as its corresponding gallery image. The performance of our method and the baseline methods on this test set is show in Table 3. We choose hard mining methods as the comparison baseline, which is a classic and state-of-the-art strategy to boost model generalization ability on hard cases. We observe that our method outperforms baseline and hard sample mining method, which verifies that the learned sampling policy help ReID model to generalize better on samples with high inner-class variation, leading to a better performance on test set with the same distribution.

To further demonstrate how our method help model learn better on hard samples, Figure 4 visualize the images with high probability to be selected by sampler at different training epoch. We observe that the sampling policy changes as the learning process progresses. At the early stage of the training (i.e. epoch 10), the high probability samples are high-resolution, clean images without strong pose variation and obvious occlusion. We believe this is because at the early stage of training when the model is not sufficiently trained, data sampler tends to select easier samples to help the model quickly converge. On the other hand, at the late stage of the training (epoch 100 / 200), we observe that more and more hard samples are selected (highlighted with red rectangle). We believe this is because when the model is almost converged, it has already sufficiently trained on the easy samples, so the data sampler should lower the importance of the easy samples to avoid over-fitting on them, and

% Label Switched	Method	Market-1501		
		mAP	Rank-1	Rank-5
5%	Baseline	75.75	89.85	96.64
	PENCIL	71.77	87.50	95.40
	<b>This work</b>	<b>79.65</b>	<b>92.22</b>	<b>97.12</b>
10%	Baseline	64.72	83.88	94.24
	PENCIL	68.24	86.16	94.92
	<b>This work</b>	<b>71.05</b>	<b>88.24</b>	<b>97.12</b>
15%	Baseline	59.54	80.14	93.35
	PENCIL	64.92	83.37	93.91
	<b>This work</b>	<b>62.61</b>	<b>83.19</b>	<b>94.15</b>

Table 4: Performance (%) comparisons of the noise resist method and our method on Market-1501 with different level of label noise added.

Method	Accuracy
Random Sampling (2018)	93.26
Focal Loss	92.59
Self-paced learning (2018)	94.31
Meta-weight Net (2019)	92.4
<b>Our method</b>	<b>96.84</b>

Table 5: Performance Comparison (%) of baseline method and our method on CIFAR10.

pay more attention to the under-fit hard samples.

**Performance on Noisy Dataset** We evaluate the influence of noisy labels on our method. Table 4 shows the performance of baseline methods and our method on the noisy dataset with different percentage of noisy label added. We choose PENCIL (Yi and Wu 2019), which is one of the state-of-the-art noise resist methods. The experiments results show that our method can resist certain level of noise in the training set, which is consistent with the real-world ReID data distribution. This is because when the level of noisiness is low, the data sampler is robust enough to select clean samples to help ReID model converge, but when the noisiness level is high, it may make ReID model fit over-fit on both clean samples and the noisy samples.

**Performance on other Task** To demonstrate our method’s effectiveness on the other task, we compare our method with existing re-sampling methods on classification task with CIFAR10 dataset (Krizhevsky and Hinton 2009). Table 5 shows the accuracy comparison of our methods and some state-of-the-art re-sampling methods and our method achieves the highest performance.

## Conclusion

This paper proposes an one-for-more objective to learn a data sampler that selects generalizable training sample to train model towards generalize on all data. The data sampler learning is integrated with ReID model learning, forming an end-to-end unified training framework. The experiments show that our model outperforms the state-of-the-art online re-sampling and re-weighting methods and show impressive generalization ability on different types of distributions.

## References

- Chawla, N. V.; Bowyer, K. W.; Hall, L. O.; and Kegelmeyer, W. P. 2002. SMOTE: synthetic minority over-sampling technique. *Journal of artificial intelligence research* 16: 321–357.
- Duan, Y.; Chen, L.; Lu, J.; and Zhou, J. 2019. Deep embedding learning with discriminative sampling policy. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 4964–4973.
- Freund, Y.; and Schapire, R. E. 1997. A decision-theoretic generalization of on-line learning and an application to boosting. *Journal of computer and system sciences* 55(1): 119–139.
- Guo, S.; Huang, W.; Zhang, H.; Zhuang, C.; Dong, D.; Scott, M. R.; and Huang, D. 2018. Curriculumnet: Weakly supervised learning from large-scale web images. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 135–150.
- He, H.; and Garcia, E. A. 2009. Learning from imbalanced data. *IEEE Transactions on knowledge and data engineering* 21(9): 1263–1284.
- Hermans\*, A.; Beyer\*, L.; and Leibe, B. 2017. In Defense of the Triplet Loss for Person Re-Identification. *arXiv preprint arXiv:1703.07737*.
- Krizhevsky, A.; and Hinton, G. 2009. Learning multiple layers of features from tiny images (Tech. Rep.). *University of Toronto*.
- Kumar, M. P.; Packer, B.; and Koller, D. 2010. Self-paced learning for latent variable models. In *Advances in Neural Information Processing Systems*, 1189–1197.
- Lin, T.-Y.; Goyal, P.; Girshick, R.; He, K.; and Dollár, P. 2017. Focal loss for dense object detection. In *Proceedings of the IEEE international conference on computer vision*, 2980–2988.
- Liu, Z.; Miao, Z.; Zhan, X.; Wang, J.; Gong, B.; and Yu, S. X. 2019. Large-Scale Long-Tailed Recognition in an Open World. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2537–2546.
- Luo, H.; Gu, Y.; Liao, X.; Lai, S.; and Jiang, W. 2019. Bag of Tricks and a Strong Baseline for Deep Person Re-Identification. In *CVPRW*.
- Ren, M.; Zeng, W.; Yang, B.; and Urtasun, R. 2018. Learning to Reweight Examples for Robust Deep Learning. In *Proceedings of the 35th International Conference on Machine Learning*, volume 80 of *Proceedings of Machine Learning Research*, 4334–4343. Stockholm: PMLR.
- Schroff, F.; Kalenichenko, D.; and Philbin, J. 2015. Facenet: A unified embedding for face recognition and clustering. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 815–823.
- Shrivastava, A.; Gupta, A.; and Girshick, R. 2016. Training region-based object detectors with online hard example mining. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 761–769.
- Shu, J.; Xie, Q.; Yi, L.; Zhao, Q.; Zhou, S.; Xu, Z.; and Meng, D. 2019. Meta-Weight-Net: Learning an Explicit Mapping For Sample Weighting. In *NeurIPS*.
- Srivastava, N.; Hinton, G.; Krizhevsky, A.; Sutskever, I.; and Salakhutdinov, R. 2014. Dropout: A Simple Way to Prevent Neural Networks from Overfitting. *Journal of Machine Learning Research* 15: 1929–1958. URL <http://jmlr.org/papers/v15/srivastava14a.html>.
- Sun, H.; Chen, Z.; Yan, S.; and Xu, L. 2019. MVP Matching: A Maximum-Value Perfect Matching for Mining Hard Samples, With Application to Person Re-Identification. In *Proceedings of the IEEE International Conference on Computer Vision*, 6737–6747.
- Sun, Y.; Zheng, L.; Yang, Y.; Tian, Q.; and Wang, S. 2018. Beyond Part Models: Person Retrieval with Refined Part Pooling. In *ECCV*.
- Tanaka, D.; Ikami, D.; Yamasaki, T.; and Aizawa, K. 2018. Joint optimization framework for learning with noisy labels. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 5552–5560.
- Thangarasa, V.; and Taylor, G. W. 2018. Self-paced learning with adaptive deep visual embeddings. *arXiv preprint arXiv:1807.09200*.
- Thulasidasan, S.; Bhattacharya, T.; Bilmes, J.; Chennupati, G.; and Mohd-Yusof, J. 2019. Combating Label Noise in Deep Learning Using Abstention. *arXiv preprint arXiv:1905.10964*.
- Wang, C.; Zhang, Q.; Huang, C.; Liu, W.; and Wang, X. 2018a. Mancs: A multi-task attentional network with curriculum sampling for person re-identification. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 365–381.
- Wang, G.; Yuan, Y.; Chen, X.; Li, J.; and Zhou, X. 2018b. Learning discriminative features with multiple granularities for person re-identification. In *ACM MM*, 274–282. ACM.
- Wu, Y.; Chen, H.; Li, S.; Gao, C.; Zhi, H.; Jiang, Y.; and Wang, Y. 2018. Person Re-Identification Based on Data Prior Distribution. In *2018 International Conference on Advanced Control, Automation and Artificial Intelligence (ACAAI 2018)*. Atlantis Press.
- Yi, K.; and Wu, J. 2019. Probabilistic End-To-End Noise Correction for Learning With Noisy Labels. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 7017–7025.
- Yu, T.; Li, D.; Yang, Y.; Hospedales, T. M.; and Xiang, T. 2019. Robust Person Re-Identification by Modelling Feature Uncertainty. In *The IEEE International Conference on Computer Vision (ICCV)*.
- Zhang, H.; Cisse, M.; Dauphin, Y. N.; and Lopez-Paz, D. 2018. mixup: Beyond Empirical Risk Minimization. In *International Conference on Learning Representations*. URL <https://openreview.net/forum?id=r1Ddp1-Rb>.
- Zhang, X.; Fang, Z.; Wen, Y.; Li, Z.; and Qiao, Y. 2017. Range loss for deep face recognition with long-tailed train-



ing data. In *Proceedings of the IEEE International Conference on Computer Vision*, 5409–5418.

Zhao, L.; Li, X.; Zhuang, Y.; and Wang, J. 2017. Deeply-Learned Part-Aligned Representations for Person Re-Identification. In *The IEEE International Conference on Computer Vision*.

Zheng, L.; Shen, L.; Tian, L.; Wang, S.; Wang, J.; and Tian, Q. 2015. Scalable Person Re-identification: A Benchmark. In *The IEEE Conference on Computer Vision and Pattern Recognition*.

Zhong, Y.; Deng, W.; Wang, M.; Hu, J.; Peng, J.; Tao, X.; and Huang, Y. 2019. Unequal-training for deep face recognition with long-tailed noisy data. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 7812–7821.

Zhou, S.; Wang, J.; Meng, D.; Xin, X.; Li, Y.; Gong, Y.; and Zheng, N. 2018. Deep self-paced learning for person re-identification. *Pattern Recognition* 76: 739–751.