

# Calibrated Domain-Invariant Learning for Highly Generalizable Large Scale Re-Identification

Ye Yuan<sup>1</sup>, Wuyang Chen<sup>1</sup>, Tianlong Chen<sup>1</sup>, Yang Yang<sup>2</sup>, Zhou Ren<sup>3</sup>, Zhangyang Wang<sup>1</sup> and Gang Hua<sup>3</sup>  
 {ye.yuan, wuyang.chen, wiwjp619, atlaswang}@tamu.edu,  
 yang.yang2@walmart.com, {renzhou200622, ganghua}@gmail.com

<sup>1</sup>Department of Computer Science and Engineering, Texas A&M University

<sup>2</sup>Walmart Technology, <sup>3</sup>Wormpex AI Research

<https://github.com/TAMU-VITA/ADIN>

## Abstract

*Many real-world applications, such as city scale traffic monitoring and control, requires large scale re-identification. However, previous ReID methods often failed to address two limitations in existing ReID benchmarks, i.e., low spatiotemporal coverage and sample imbalance. Notwithstanding their demonstrated success in every single benchmark, they have difficulties in generalizing to unseen environments. As a result, these methods are less applicable in a large scale setting due to poor generalization. In seek for a highly generalizable large-scale ReID method, we present an adversarial domain-invariant feature learning framework (ADIN) that explicitly learns to separate identity-related features from challenging variations, where for the first time “free” annotations in ReID data such as video timestamp and camera index are utilized. Furthermore, we find that the imbalance of nuisance classes jeopardizes the adversarial training, and for mitigation we propose a calibrated adversarial loss that is attentive to nuisance distribution. Experiments on existing large-scale person/vehicle ReID datasets demonstrate that ADIN learns more robust and generalizable representations, as evidenced by its outstanding **direct transfer** performance across datasets, which is a criterion that can better measure the generalizability of large scale Re-ID methods.*

## 1. Introduction

The increasing usage of sensors, especially surveillance cameras, in smart communities and cities, has resulted in immense opportunities in developing large-scale computer vision and machine learning algorithms for Urban Informatics. As a core algorithmic component of many camera-based applications, the re-identification of subjects across mul-


iple cameras, known as ReID, has been one of the most demanded capabilities. For example, person ReID targets to match and return images of a probe person from a large-scale gallery set collected from different cameras, which is an important in security and surveillance. As another example, vehicle ReID uses traffic cameras as citywide sensors to optimize flows and manage traffic accidents. These ReID problems have recently drawn explosive attention from both academia and industry. Despite notable research progress, there remain to be major gaps between the research efforts and the practical needs in large-scale deployment.

### 1.1. The Generalizability Gaps: Low Coverage and Sample Imbalance

**High generalizability is essential to deploy ReID at large scale.** ReID is by nature an “open-world” problem, which is necessitated to generalize in two aspects. First, the ReID models have to generalize to subjects (*e.g.*, person or vehicle) that were unseen in the training set. Second, ReID shall also generalize to re-identifying subjects that undergo unseen variations induced by changes in background, illumination, viewpoint – *i.e.*, generalizing to unseen “combinations” of subjects and accompanied variations. We call those unwanted variations as **nuisances** since they are irrelevant to the subject identities.

With the rapidly expanding smart city/community camera networks, good generalizability becomes the *ultimate goal* for the practical deployment of large-scale ReID systems: it has to stay effective to new subjects, can scale up to new locations, and be reliable over time. However, most existing ReID algorithms may not have addressed these well, which may hinder their deployment in practice.

**Gap #1: Low Spatiotemporal Coverage.** Existing ReID datasets are limited in data volume. They are restricted in not only subject numbers but also spatial/temporal cov-



**Figure 1:** The sample imbalance issue in ReID datasets (MSMT17 [1] as an example). The imbalance manifests in: (left) histogram of image numbers per subject; (middle) histogram of camera numbers in which each subject was captured; (right) histogram of time period numbers (1/2/3: morning/noon/afternoon) during which each subject was captured. The latter two histograms show to be clearly skewed towards the lower end, implying the “localized” patterns of person activities in MSMT17. Similar observations can be found in other peer datasets.

erage, making them oversimplified in reference to the complexity and diversity of the large-scale scenes. A recent study [2] showed that in 2014, there were 125 video surveillance cameras per thousand people in the U.S.; whereas most ReID datasets were collected only from 10 or fewer cameras (see Table 1). What is worse, most datasets hand-picked video frames captured in similar outdoor environments and/or under normal lighting conditions; however, practical ReIDs need to cope with drastically diverse locations, indoor-outdoor matching, intensive day-long illumination variations, and more.

**Gap #2: Sample Imbalance<sup>1</sup>.** We believe this issue remains yet *overlooked* in the ReID community, especially when the open scenarios in real cities go beyond the controlled conditions in training datasets. First, a ReID dataset consists of images from different subjects, where every subject has an indefinite number of images (Fig.1 left), making the subject class distribution *non i.i.d.* Similarly, different nuisances may also appear in a dataset with different frequencies. What is worse, a subject may (very likely) be only captured by a small portion of cameras (Fig.1 middle) within a small portion of time periods (Fig.1 right). Intuitively, a person or a vehicle usually appears most in certain regions within certain hours, rather than being a wanderlust anywhere anytime “uniformly” in a city. Hence, the conditional distribution of nuisances given a subject is also *heavily non i.i.d.*: some subjects may display strong yet superficial correlations with some nuisances, making it highly challenging to decouple them.

Considering these gaps, we argue that evaluating (and even overfitting) on single datasets is not helpful for designing algorithms that would be practically deployed in real-life large-scale scenarios.

## 1.2. Our Solution and Contributions

The low spatiotemporal coverage coupled with the inherent sample imbalance, have placed jeopardy for practical large-scale ReID. Existing ReID algorithms trained on a single dataset are prone to overfitting nuisances of the training set, and therefore suffering from poor generalizability, as indicated by poor direct transfer performance to unseen datasets (Fig. 2 and Table 3).

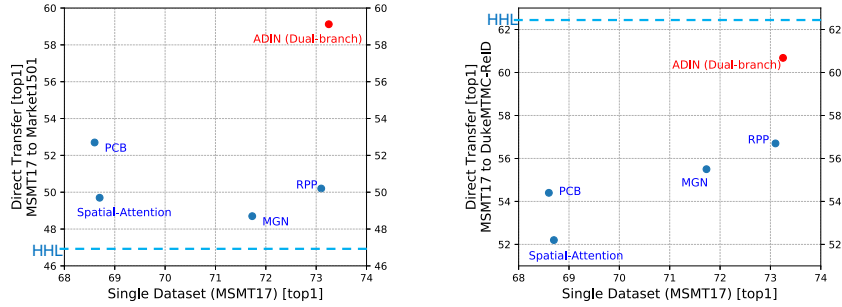
Fortunately, most scene-related nuisances are caused by *camera-specific* and/or *time-specific* factors (Fig. 1 middle and right). Since *video timestamp* or *camera index* are freely available in video surveillance as metadata and are provided by almost all existing person/vehicle ReID datasets, they can be potentially utilized as auxiliary supervision, although few image-based ReID methods have taken advantage of them.

This paper aims to improve the **generalizability** of ReID models in large-scale settings, by resorting to a novel domain-invariant feature learning perspective. Inspired by [4, 5], we consider samples (of different subjects) with the same nuisance to be from one domain (such as images captured by the same fixed camera, or in the same time period). This is because scene-related changes (background, illumination, viewpoint, etc.) heavily dominate the appearances of images. Different types of nuisances hence becomes domain-specific features. In contrast, one subject can be captured at different cameras and time periods, and the subject’s identity features should apparently remain domain-invariant. Therefore, our main idea is to extract ReID features that can: (1) be utilized to faithfully classify subjects into correct classes; (2) be resilient and invariant to those identified nuisances – in other words: no reliable classifier can be trained on those features to predict those nuisances.

In addition to setting up a new goal of generalizability,

<sup>1</sup>In ReID, “sample imbalance” traditionally refers to the negative samples being much more than positive ones during pairwise training [3]. We are using this term in an apparently different context.

<sup>2</sup>HHL uses images from both source and target domain for domain adaptation, and thus has no single-dataset performance. We use a horizontal line to represent its domain adaptation performance.



**Figure 2:** Top1 accuracy on a single-dataset (MSMT17 [1]) and direct transfer from MSMT17 to Market1501 [6] (left) and to DukeMTMC-ReID [7, 8] (right). In contrast to our ADIN (red dot in top-right) which achieved competitive performance on both a single dataset and direct transfer, we find other (single-dataset) top-performers suffer from very poor generalizability to unseen domains, indicating the misaligned goal between overfitting small-scale single dataset and generalizing to large-scale unseen scenarios in real life. See section 4.1 for details. Methods studied: Spatial-Attention [9], PCB[10], RPP [10], MGN [11], HHL<sup>2</sup>[12].

**Table 1:** Publicly available benchmarks for person/vehicle ReID: the number of cameras, identities, and average bounding boxes per identity.

	Benchmark	#cameras	#ID	#boxes/ID
person	Market-1501 [6]	6	1501	21.5
	DukeMTMC-ReID[7, 8]	8	1812	20.1
	CUHK03[13]	2	1467	9.0
	<b>MSMT17[1]</b>	<b>15</b>	<b>4101</b>	<b>30.8</b>
vehicle	VeRi[14]	20	619	64.6
	VehicleID[15]	2	26267	8.4
	<b>VeRi-776[16]</b>	<b>20</b>	<b>776</b>	<b>63.6</b>

our contributions include a corresponding new algorithm as well as a new evaluation criterion. We formulate our adversarial domain-invariant learning framework (**ADIN**), by taking advantage of “free” annotations like video timestamp and camera index, to separate identity-related features from scene-specific nuisances. To our best knowledge, we are the first to utilize those “free” annotations for image-based ReID, to effectively suppress the overfitting of nuisances. Moreover, we find the imbalance of nuisance distribution w.r.t. subjects hampers the adversarial learning. A novel calibrated adversarial loss is therefore introduced to tackle the nuisance class imbalance for ADIN. Measured by a new **direct transfer** performance criterion on popular large-scale ReID benchmarks, our ADIN demonstrates outstanding generalizability and outperform previously reported results and even some that rely on domain adaptation using target data (Fig.2).

## 2. Related Work

### 2.1. ReID Datasets

The disconnection between research-level datasets and community/city-level video warehouse remains to hinder the real-life applications of ReID. Table 1 summarizes mainstream person/vehicle ReID datasets. For person ReID, considering that even in a grocery store there are usually dozens

even more than 100 cameras [17] and over 550 visitors per day [18], current datasets are more or less overly simplistic. More specifically, the Market-1501 [6], DukeMTMC-ReID [7, 8] and CUHK03 [13] are all collected in small outdoor regions, and in short time periods (usually well-lighted daytime). The latest MSMT17 dataset [1] led positive progress towards real large-scale usage, by including geo-spatially diverse cameras (both indoor and outdoor) and varying time periods (morning, noon and afternoon) and illuminations.

Vehicle ReID witness similar situations, where exiting benchmarks’ scale and diversity are still far from being comparable to reality. Previous datasets such as VehicleID [15] have small camera or vehicle numbers, as well as limited viewpoints. A recent VeRi-776 dataset [16] presented a relatively realistic benchmark with cameras spanning a large spatial coverage and other variations, which is one-step close to being representative for large-scale vehicle ReID.

### 2.2. ReID Evaluation Metrics

The standard ReID pipeline picks a dataset, learning the model from its training set and evaluating the model’s retrieval accuracy or mean average prevision (mAP) on the held out testing set (with non-overlapping subjects). However, this **single-dataset evaluation is often insufficient in reflecting true generalizability** (Fig.2) since they overlook a fact, *i.e.*, due to the low coverage of most datasets, the training and testing sets of the same ReID dataset tend to be highly similar in terms of spatiotemporal nuisances (even overlapping or sharing camera IDs). Therefore, a high accuracy/mAP on the same testing set may be misleading, as that could be a result of nuisance overfitting.

In view of the above, increasing attention has been paid to domain adaption in ReID recently, *i.e.*, training on one source dataset, tuning the trained model on some different target domain data, and finally evaluating model accuracy/mAP on the target dataset. Domain adaptation methods [12] emphasize the generalizability of ReID to new data. Unfortunately,

they require target domain data (sometimes even auxiliary attribute annotations in target domain [19]) for re-training purposes. Considering the city growth as well as the explosive increase of cameras, it is unrealistic to collect new data and re-train ReID models for every new domain (e.g., a new camera or a group of cameras in a local region), making it non-trivial for domain adaptation to scale up.

In contrast, we advocate a far more challenging but practically evaluation criterion: **direct transfer** performance across datasets, to measure ReID model generalizability (see section 4.1).

### 2.3. Algorithms for Improving ReID Generalizability

**Data Augmentation for ReID.** Tian et al. [20] proposed to generate images of the same identity with different random backgrounds. In Ma et al. [21] a camera-invariant descriptor subspace is learned and the camera styles are transferred to each sample. [22] first learned a set of disentangled foreground, background and pose factors, then re-composed them into novel samples. Data augmentation will amplify the training burden. Moreover, they still suffer when transferred to an unseen dataset, since no single training dataset (as being low-coverage itself) can cover sufficient real-world variations.

**Domain Adaptation for ReID.** Deng et al. [23] proposed to generate a new training set of images whose identities were from the labeled source domain, while the camera styles were translated from the unlabeled target domain. Zhong et al. [12] introduced a Hetero-Homogeneous Learning (HHL) method to learn person embedding with camera variances and domain connectedness, through inter- and intra-domain pairwise contrastive learning. An unsupervised image translation approach was presented in [24] for source models to learn the style of the target domain. Different from them, our ADIN is “directly transferable” to unseen domains, without seeing or re-training on target data.

## 3. The Proposed Approach: ADIN

### 3.1. Problem Formulation

Given a training image  $X$  with the identity labels  $Y_I$  and the (freely) available nuisances label  $Y_N$  (one or multiple, such as camera ID, video timestamp, etc.), our goal is to learn a feature representation  $f_E(X)$  that is highly **relevant** to the identity label, yet being invariant or **irrelevant** to the nuisances label. Using a function  $R$  to represent the correlation between the feature and the label, our learning goal is mathematically described as:

$$R(f_E(X), Y_I) \approx R(X, Y_I), \quad R(f_E(X), Y_N) \ll R(X, Y_N). \quad (1)$$

We adopted an identity prediction module  $f_I$  which projects the feature  $f_E(X)$  into identity-related features, and a nuisance prediction module  $f_N$  that extracts scene-specific nuisances from  $f_E(X)$ . Without loss of generality, both of them are assumed to have softmax-form outputs. Note that  $f_E$ ,  $f_I$  and  $f_N$  all need to be learned together. Their interactions provide mutual supervision. In particular,  $f_N$  will serve as an “adversary” role.

To evaluate  $R$  practically, a straightforward choice is to use two standard classification-oriented loss functions  $L_I$  and  $L_N$  (e.g., cross-entropy) for  $f_I$  and  $f_N$  respectively and minimize the classification error rate of  $Y_I$  from  $f_E(X)$ , while maximizing the classification error rate of  $Y_N$  from  $f_E(X)$ . Our task then becomes to simultaneously train  $f_E$ ,  $f_I$  and  $f_N$ , so as to minimize  $L_I(f_I(f_E(X)), Y_I)$  meanwhile maximizing  $L_N(f_N(f_E(X)), Y_N)$ .

$$\min_{f_E} L_I(f_I(f_E(X)), Y_I), \quad \max_{f_E} L_N(f_N(f_E(X)), Y_N). \quad (2)$$

Maximizing  $L_N(f_N(f_E(X)), Y_N)$  is not straightforward to implement. Previous work [25] reversed the sign of gradient computed from minimizing  $L_N(f_N(f_E(X)), Y_N)$ , i.e., using gradient ascent. However, we observed in experiments that the reverse gradient approach yielded unstable training process. Instead, we introduce a new  $L_{adv}$  loss, to encourage the *disparity* between  $f_N(f_E(X))$  and  $Y_N$ : a smaller  $L_{adv}$  value is expected to indicate a *worse* correlation between them. A detailed discussion about the choice of  $L_{adv}$  will be presented in section 3.2.

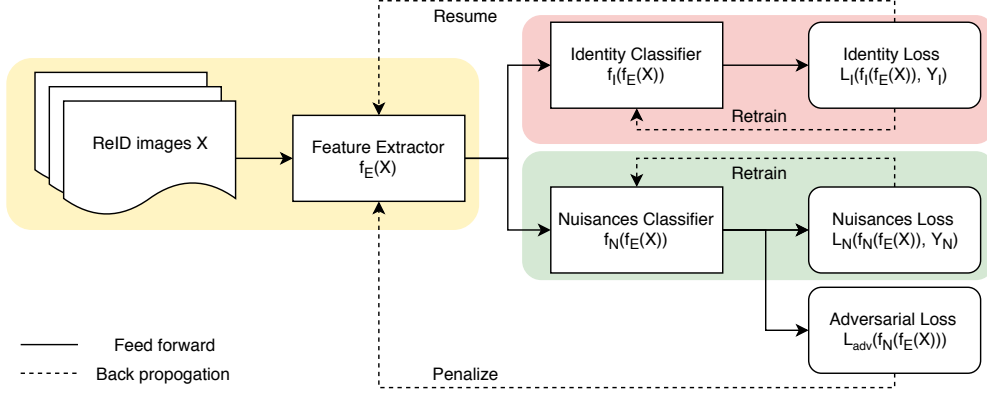
Finally, the training goal of ADIN is represented below ( $\beta > 0$  is a scalar):

$$\min_{f_E} L_I(f_I(f_E(X)), Y_I) + \beta L_{adv}(f_N(f_E(X)), Y_N). \quad (3)$$

Meanwhile, in order to keep adversarial domain-invariant feature learning effective so as to learn meaningful  $f_E$ , we need to also maintain  $f_N$  to be a strong competitor. That implies a *hidden constraint*, i.e., avoiding  $L_N(f_N(f_E(X)), Y_N)$  growing large too quickly, in which case  $f_N$  becomes to have too poor nuisance classification ability so that it cannot make a useful adversary.

### 3.2. Calibrated Adversarial Loss for Imbalanced Nuisances

As noted in section 1.1, both subject and nuisance classes (conditioned on the subject) suffer from sample imbalances. We experimentally observed the subject imbalance to have less severe impact on ReID performance (e.g., comparing using standard and reweighted softmax loss), and therefore keep using a standard softmax function for  $L_I$ . However, the nuisance class imbalance was found to cause considerable training instability and performance degradation for the



**Figure 3:** Overview of the ADIN framework, and illustration of its training strategy. See section 3.3 for details.

adversarial learning. We thus focus a detailed discussion on how we derive a robust  $L_{adv}$  for the imbalanced nuisances.

We denote  $c = [c_1, \dots, c_K]$  as the softmax-form output of  $f_N$ , where  $K$  is the total nuisance class number. We next present three options that we tried for  $L_{adv}$ , among which our proposed new Option #3 is experimentally validated to be the best choice for ADIN (see section 4.3 for details).

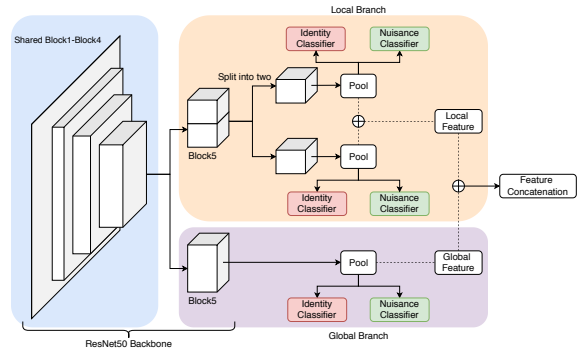
**Option #1: Reverse Gradient (RG).** One possibility is to adopt the reversal gradient layer [26]. It computes the gradient for minimizing the cross-entropy between  $f_N(f_E(X))$  and  $Y_N$ , then reversing the gradient sign. However, this objective becomes problematic in our case, as it was observed to cause large fluctuations in the training curve and failure of convergence. Moreover, when both  $f_N$  and  $f_E$  are initialized from pre-trained models (practically improving convergence and results), the gradients start with very small magnitudes and the model updates become too slow. RG is written as ( $Y^*$  is the true label):

$$L_{adv}(X, Y_N) = -L_N = \sum_{k=1}^K \mathbb{1}_{[k=Y^*]} \log(c_k) \quad (4)$$

**Option #2: Negative Entropy (NE).** An alternative is to minimize the negative entropy function of the softmax vector (or equivalently, its *cross-entropy* with uniform distribution), as to encourage “uncertain” predictions of nuisance attributes (e.g., camera ID and video timestamps) from the extracted ReID features. The rationale is that, if the nuisance prediction is only as good as the random guess (uniform distribution over all classes), then the feature is not informed of nuisances and therefore can generalize to unseen nuisances. NE could be written as

$$L_{adv}(X, Y_N) = \sum_{k=1}^K c_k \log(c_k). \quad (5)$$

Importantly, although  $Y_N$  does not explicitly occur in the loss form, it will still be utilized in re-training  $f_N$  to make a



**Figure 4:** Overview of the dual-branch backbone.

sufficiently strong competitor (section 3.3). We previously also tried the KL Divergence and the Jensen-Shannon Divergence between the softmax and uniform distribution, but NE appears to work best in practice.

**Option #3 (Proposed): Calibrated Negative Entropy Loss (CaNE).** Despite boosting uncertainty, NE overlooks the practical imbalance of nuisance class distribution w.r.t. subjects. A well-known solution is to add a modulating factor to cross-entropy loss, ensuring that the majority class/easy decisions do not overwhelm the loss [27]. We propose a reweighted form of NS, called Calibrated Negative Entropy Loss (CaNE), to make  $L_{adv}$  attentive to the skewed nuisance distribution

$$L_{adv}(X, Y_N) = \sum_{k=1}^K p_k c_k \log(c_k), \quad (6)$$

where  $p_k$  denotes the nuisance class distribution in the given training set. To our best knowledge, there has been no similar discussion addressing the class imbalance issue in (adversarial) domain adaption among existing ReID works.

### 3.3. Training Strategy Overview

Figure 3 overviews the concrete training workflow of ADIN, which consists of three modules: feature extractor  $f_E$ ,

**Table 2:** Ablation study of  $L_{adv}$  (direct transfer DukeMTMC-ReID  $\rightarrow$  Market1501).

Settings		DukeMTMC-ReID $\rightarrow$ Market1501			
		top1	top5	top10	mAP
ResNet50 (baseline)		46.8	63.5	70.3	19.0
ADIN	ResNet50 + Reverse Gradient	Unable to converge			
	ResNet50 + NE	48.8	66.2	72.7	20.4
	ResNet50 + CaNE	51.7	68.6	76.0	22.1
Dual-branch (baseline)		54.8	71.7	77.6	25.9
ADIN	Dual-branch + Reverse Gradient	Unable to converge			
	Dual-branch + NE	55.9	72.5	78.6	26.5
	<b>Dual-branch + CaNE</b>	<b>57.2</b>	<b>73.0</b>	<b>80.0</b>	<b>27.4</b>

subject identity classifier  $f_I$ , and nuisance classifier  $f_N$ .  $f_E$  takes the image  $X$  as input and outputs the feature  $f_E(X)$ , which is then passed through  $f_I$  and  $f_N$ . Both  $f_I$  and  $f_N$  aim to accurately predict their corresponding labels from the learned features. The training of  $f_E$  strives to boost the prediction of  $f_I(f_E(X))$ , while suppressing the prediction of  $f_N(f_E(X))$ . It is important to keep  $f_N$  strong to maintain a meaningful competition for learning nontrivial  $f_E$ .

In practice, we implement the training using an iterative strategy. We initialize  $f_E$ ,  $f_I$  and  $f_N$  by jointly training the feature extractor  $f_E$  and identity classifier  $f_I$ , and then fixing  $f_E$  and pre-training  $f_N$  solely on top of that. Afterwards, we alternate between optimizing two sub-problems:

$$\min_{f_E, f_N} L_{adv}(f_N(f_E(X))), \quad \min_{f_E, f_I} L_I(f_I(f_E(X)), Y_I). \quad (7)$$

In each alternating round, we optimize the first objective until the validation error of identity classification reducing below a pre-set  $threshold_{I-target}$ . We then switch to optimizing the second objective, meanwhile monitoring the resulting changes on the identity classification validation error (since  $f_E$  is altered): if it drops below another pre-set  $threshold_{I-trigger}$ , we will switch back to the first object and start the next round of alternations. See supplementary for full training details.

## 4. Experiment

### 4.1. Direct Transfer across Datasets as ReID Evaluation Metric

In view of current ReID dataset limitations, we propose to choose a new evaluation metric to reflect our ultimate goal of generalizability. In real large-scale ReID applications, one would expect a trained model to “automatically” scale up to as many unseen, different scenarios as possible. Since different existing ReID datasets were collected in very diverse settings, we propose to use the **direct transfer** performance, by directly applying a ReID model trained on one dataset’s training set onto another’s testing set to measure accuracy/mAP, **without any re-training or adaptation**. The direct transfer performance explicitly takes into

account “zero-shot” generalizability and penalizes the overfitting of setting/scene-specific nuisances. We argue that it should be as important a performance indicator for ReID as traditional accuracy/mAP on the same dataset, if not more.

### 4.2. Datasets and Model Implementation

We first perform “small-scale”, proof-of-concept ablation studies using DukeMTMC-ReID [7, 8] and Market-1501 [6] (section 2), then go to two large-scale datasets: MSMT17 [1] for person ReID, and VeRi-776 [16] for vehicle ReID (section 4.1). For all datasets, we use target subject IDs (person or vehicle) as  $Y_I$ . For  $Y_N$ , we take: (i) camera IDs on DukeMTMC-ReID and on VeRi-776; and (ii) both camera IDs and video timestamps on MSMT17, through a multi-label adversarial classifier.

As a general framework, ADIN can take any backbone for  $f_E$ ,  $f_I$  and  $f_N$ . In section 4.3, we first test our ADIN with  $f_E$  being a basic ResNet50 [28] to illustrate the effectiveness of our adversarial training. Afterwards, we adopt a more sophisticated dual-branch feature extractor for  $f_E$ , as inspired by [10, 11, 29], to demonstrate further boosted performance over state-of-the-arts. The configuration of the dual-branch model is depicted in Fig.4. On top of the  $f_E$ , we append two simple classifiers as  $f_I$  and  $f_N$  (Fig. 3), either taking two fully connected layers.  $L_I$  is always implemented using the hybrid loss of cross-entropy and center loss [30]. An ablation study of  $L_{adv}$  is presented in section 3.2; after that, the Calibrated Negative Entropy (CaNE) loss will be our default  $L_{adv}$  unless otherwise specified.

More details of our models can be found in supplementary. All codes and pre-trained models will be released upon acceptance.

### 4.3. Ablation study of the adversarial loss $L_{adv}$

Table 2 displays a step-by-step comparison for choosing  $L_{adv}$ , with the direct transfer performance from DukeMTMC-ReID (source domain) to Market1501 (target domain) as the indicator. Without the adversarial domain-invariant training, both the ResNet50 and Dual-branch backbones achieve low direct transfer accuracy, due to the domain discrepancy across two datasets. We also empirically observe the adversarial effect provided by the reverse gradient (RG) hard to converge, owing to the gradient vanishing/explosion and its sensitivity to the loss magnitude from the nuisance classifier  $f_N$ . With the negative entropy (NE) loss, our adversarial domain-invariant training forces the entropy of the nuisance classifier’s prediction to be maximized, leading to nuisance-uninformative features learned by the feature extractor  $f_E$  and reliable direct transfer performance. More importantly, as pointed out in section 1.1, the sampling of nuisances are imbalanced, which intrinsically results in imbalanced levels of adversarial effects on each nuisance. Thus our proposed calibrated negative loss (CaNE) further en-



**Table 3:** Direct transfer performance from MSMT17 [1] to DukeMTMC-ReID [7, 8] and to Market1501 [6]. \* indicates method using images from both source and target domain.

	MSMT17 → DukeMTMC-ReID			MSMT17 → Market1501		
	top1	top5	mAP	top1	top5	mAP
Spatial-Attention[9]	52.2	68.1	32.9	49.7	68.9	25.1
PCB[10]	54.4	69.6	34.6	52.7	71.3	26.7
RPP[10]	56.7	71.4	36.7	50.2	70.7	26.3
MGN[11]	55.5	70.2	35.1	48.7	66.9	25.1
<b>HHL[12]*</b>	<b>62.2</b>	<b>78.8</b>	31.4	46.9	61.0	27.2
ResNet50 (baseline)	49.7	65.7	28.2	47.7	64.3	21.2
ResNet50 + CaNE	52.6	67.9	30.4	50.1	66.4	22.5
Dual-branch	59.5	73.5	38.4	57.8	73.9	29.4
<b>Dual-branch + CaNE</b>	<b>60.7</b>	<b>74.7</b>	<b>39.1</b>	<b>59.1</b>	<b>75.4</b>	<b>30.3</b>

**Table 4:** Direct transfer performance from VeRi-776 [14] to VehicleID[15]. \* indicates method using images from both source and target domain.

Method	Test size = 1600			Test size = 3200		
	top1	top5	mAP	top1	top5	mAP
RAM[31]	30.5	49.5	39.5	24.5	40.3	32.4
Spatial-Attention[9]	39.5	57.2	47.9	33.7	49.6	41.6
PCB[10]	41.3	58.8	49.7	35.4	51.4	43.2
RPP[10]	40.6	58.4	49.1	35.0	51.1	42.9
MGN[11]	39.9	62.4	50.6	32.7	53.1	42.7
<b>DAVR[24]*</b>	<b>45.2</b>	<b>64.0</b>	<b>49.7</b>	<b>38.7</b>	<b>55.9</b>	<b>42.9</b>
ResNet50 (baseline)	42.3	58.5	46.2	36.1	52.2	39.9
ResNet50 + CaNE	43.3	59.7	47.2	37.0	53.4	40.9
Dual-branch	47.3	65.3	51.6	41.2	57.9	45.3
<b>Dual-branch + CaNE</b>	<b>48.7</b>	<b>67.3</b>	<b>53.1</b>	<b>42.1</b>	<b>59.5</b>	<b>46.3</b>

ables the adversarial training to be attentive w.r.t. different nuisances frequencies. Table 2 shows that both backbones benefit most from our proposed CaNE adversarial loss.

It is worth noting that even trained within a small-scale domain like DukeMTMC-ReID, the generalizability of both of the two backbones can be boosted by ADIN.

#### 4.4. Direct Transfer between Datasets without Re-training or Adaption

We evaluate three direct transfer cases, two on person ReID: MSMT17 → DukeMTMC-ReID, MSMT-17 → Market1501; and one on vehicle ReID: VeRi-776 [14] → VehicleID [15]. As comparison baselines, we train the same dual-branch backbones (without any adversarial learning) on the source datasets, and test their direct transfer performance too. We train and compare with several state-of-the-art ReID models on MSMT17: Spatial-Attention [9], PCB [10], RPP [10], MGN [11] (Person ReID); and RAM [31] (Vehicle

ReID). We further compare with existing best performers of domain adaptation: HHL [12] (Person ReID), and DAVR [24] (Vehicle ReID), which reported the current best transfer results between DukeMTMC-ReID/Market1501, and from VeRi-776 [14] to VehicleID [15], respectively. Note that both HHL and DAVR need to use (unlabeled) target domain data and perform extra (re-)training for the source domain models, while ours need not: the comparisons are thus apparently to our competitors' advantage.

As can be seen from Tables 3 and 4, while baselines without adversarial learning fail to transfer well as expected, ADIN demonstrates highly impressive results on all three transfer cases. In particular, by training on MSMT17 and directly transferring, ADIN not only surpasses the direct transfer results from other methods but also outperforms state-of-the-art ReID domain adaption model HHL [12] and DAVR [24], while costing literally none of their hassles such as (re-)training.



**Figure 5:** Retrieval results on the DukeMTMC-ReID (a), MSMT17 (b, c), and VeRi-776 (d, e). The leftmost image in each panel is a query image. For the five columns in each panel, the top row shows top-5 retrieval results using a vanilla ResNet50 model [28], and the bottom row shows top-5 results using ResNet50 adopted with ADIN framework. Green boxes mark the correct matches, while red boxes denote the wrong matches. The vanilla ResNet50 tend to retrieve images with similar nuisances (illumination, viewpoints, scene backgrounds, etc.), while ADIN successfully eliminates nuisances and leads to more robust matching under drastic visual appearance changes.

In contrast to our ADIN, we find other (single-dataset) top-performers [9, 10, 11] generalize very poorly to unseen domains, indicating the misaligned goals between overfitting small-scale single dataset, and generalizing to large-scale unseen scenarios in real life. As in Fig.2, our ADIN framework lies in the top-right corner, while others stay in the left region with high single-dataset accuracy but poor direct transfer performance. We believe the effective direct transfer is the right choice for evaluating and promoting larger-scale ReID practice, and hope our proposals and arguments could invoke more discussions in the community. We include the full detailed results in the supplementary.

Figure 5 shows five visual retrieval examples. In both queries, the spatiotemporal nuisances (*e.g.* the door of same geo-location, certain lighting condition or glare) have a strong presence in images. As can be seen in the top row of each case, the baselines overfit background, tending to retrieve images with similar nuisances (illumination, viewpoints, scene backgrounds, etc.). In contrast, ADIN successfully eliminates them, and leads to much more robust matching under drastic visual appearance changes, as seen in the bottom rows.

In sum, ADIN proves to make a substantial improvement

in overcoming the generalizability challenge: the ADIN feature extractor learned on one dataset is directly generalizable to others, without seeing or adapting on any new data. To our best knowledge, ADIN is the first CNN-based ReID model that can establish strong direct transfer performance. It sets up state-of-the-art generalizability for ReID, which we believe is valuable for pursuing real-world large-scale ReID applications.

## 5. Conclusion

This paper proposes the adversarial domain-invariant (**ADIN**) learning framework, which remarkably improves the generalizability of ReID models and resolves the nuisances-overfitting problem. Free annotations like video timestamp and camera index are for the first time utilized. In extensive experiments, measured by the new **direct transfer** performance criterion, ADIN exhibits impressive generalization to unseen datasets without any fine-tuning or re-training. We hope to draw more research interests to address the ReID generalizability and to pursue more effective direct transfer.



## References

- [1] Longhui Wei, Shiliang Zhang, Wen Gao, and Qi Tian. Person transfer gan to bridge domain gap for person re-identification. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2018.
- [2] Number of video surveillance cameras per thousand people in 2014, by country. <https://www.statista.com/statistics/484956/number-of-surveillance-cameras-per-thousand-people-by-country/>, 2015.
- [3] Yicheng Wang, Zhenzhong Chen, Feng Wu, and Gang Wang. Person re-identification with cascaded pairwise convolutions. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1470–1478, 2018.
- [4] Zhenyu Wu, Zhangyang Wang, Zhaowen Wang, and Hailin Jin. Towards privacy-preserving visual recognition via adversarial training: A pilot study. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 606–624, 2018.
- [5] Haotao Wang, Zhenyu Wu, Zhangyang Wang, Zhaowen Wang, and Hailin Jin. Privacy-preserving deep visual recognition: An adversarial learning framework and a new dataset. *arXiv preprint arXiv:1906.05675*, 2019.
- [6] Liang Zheng, Liyue Shen, Lu Tian, Shengjin Wang, Jingdong Wang, and Qi Tian. Scalable person re-identification: A benchmark. In *The IEEE International Conference on Computer Vision (ICCV)*, December 2015.
- [7] Ergys Ristani, Francesco Solera, Roger Zou, Rita Cucchiara, and Carlo Tomasi. Performance measures and a data set for multi-target, multi-camera tracking. In *The European Conference on Computer Vision (ECCV)*, September 2016.
- [8] Zhedong Zheng, Liang Zheng, and Yi Yang. Unlabeled samples generated by gan improve the person re-identification baseline in vitro. In *The IEEE International Conference on Computer Vision (ICCV)*, Oct 2017.
- [9] Haoran Wang, Yue Fan, Zexin Wang, Licheng Jiao, and Bernt Schiele. Parameter-free spatial attention network for person re-identification. *arXiv preprint arXiv:1811.12150*, 2018.
- [10] Yifan Sun, Liang Zheng, Yi Yang, Qi Tian, and Shengjin Wang. Beyond part models: Person retrieval with refined part pooling (and a strong convolutional baseline). In *The European Conference on Computer Vision (ECCV)*, September 2018.
- [11] Guanshuo Wang, Yufeng Yuan, Xiong Chen, Jiwei Li, and Xi Zhou. Learning discriminative features with multiple granularities for person re-identification. In *2018 ACM Multimedia Conference on Multimedia Conference*, pages 274–282. ACM, 2018.
- [12] Zhun Zhong, Liang Zheng, Shaozi Li, and Yi Yang. Generalizing a person retrieval model hetero- and homogeneously. In *The European Conference on Computer Vision (ECCV)*, September 2018.
- [13] Wei Li, Rui Zhao, Tong Xiao, and Xiaogang Wang. Deepreid: Deep filter pairing neural network for person re-identification. In *CVPR*, 2014.
- [14] Xinchen Liu, Wu Liu, Huadong Ma, and Huiyuan Fu. Large-scale vehicle re-identification in urban surveillance videos. In *2016 IEEE International Conference on Multimedia and Expo (ICME)*, pages 1–6. IEEE, 2016.
- [15] Hongye Liu, Yonghong Tian, Yaowei Yang, Lu Pang, and Tiejun Huang. Deep relative distance learning: Tell the difference between similar vehicles. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2167–2175, 2016.
- [16] Xinchen Liu, Wu Liu, Tao Mei, and Huadong Ma. A deep learning-based approach to progressive vehicle re-identification for urban surveillance. In *European Conference on Computer Vision*, pages 869–884. Springer, 2016.
- [17] Amazon go debuts, and its prying cameras foil our shoplifting attempts. <https://arstechnica.com/information-technology/2018/01/we-test-the-worlds-first-amazon-go-watch-you-shop-grocery-store/>, 2018.
- [18] Analysts: Amazon go stores bring in 50% more revenue than typical c-stores. <https://csnews.com/analysts-amazon-go-stores-bring-50-more-revenue-typical-c-stores>, 2019.
- [19] Jingya Wang, Xiatian Zhu, Shaogang Gong, and Wei Li. Transferable joint attribute-identity deep learning for unsupervised person re-identification. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2018.
- [20] Maoqing Tian, Shuai Yi, Hongsheng Li, Shihua Li, Xuesen Zhang, Jianping Shi, Junjie Yan, and Xiaogang Wang. Eliminating background-bias for robust person re-identification. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2018.
- [21] Zhun Zhong, Liang Zheng, Zhedong Zheng, Shaozi Li, and Yi Yang. Camera style adaptation for person re-identification. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2018.
- [22] Liqian Ma, Qianru Sun, Stamatios Georgoulis, Luc Van Gool, Bernt Schiele, and Mario Fritz. Disentangled person image generation. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2018.
- [23] Weijian Deng, Liang Zheng, Qixiang Ye, Guoliang Kang, Yi Yang, and Jianbin Jiao. Image-image domain adaptation with preserved self-similarity and domain-dissimilarity for person re-identification. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2018.
- [24] Jinjia Peng, Huibing Wang, and Xianping Fu. Cross domain knowledge learning with dual-branch adversarial network for vehicle re-identification. *arXiv preprint arXiv:1905.00006*, 2019.
- [25] Yaroslav Ganin and Victor Lempitsky. Unsupervised domain adaptation by backpropagation. *arXiv preprint arXiv:1409.7495*, 2014.
- [26] Yaroslav Ganin, Evgeniya Ustinova, Hana Ajakan, Pascal Germain, Hugo Larochelle, François Laviolette, Mario Marchand, and Victor Lempitsky. Domain-adversarial training of neural networks. *Journal of Machine Learning Research*, 17(59):1–35, 2016.

- [27] Tsung-Yi Lin, Priya Goyal, Ross Girshick, Kaiming He, and Piotr Dollár. Focal loss for dense object detection. In *Proceedings of the IEEE international conference on computer vision*, pages 2980–2988, 2017.
- [28] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016.
- [29] Tianlong Chen, Shaojin Ding, Jingyi Xie, Ye Yuan, Wuyang Chen, Yang Yang, Zhou Ren, and Zhangyang Wang. Abd-net: Attentive but diverse person re-identification. *ICCV*, 2019.
- [30] Yandong Wen, Kaipeng Zhang, Zhifeng Li, and Yu Qiao. A discriminative feature learning approach for deep face recognition. In *The European Conference on Computer Vision (ECCV)*, pages 499–515. Springer, 2016.
- [31] Xiaobin Liu, Shiliang Zhang, Qingming Huang, and Wen Gao. Ram: a region-aware deep model for vehicle re-identification. In *2018 IEEE International Conference on Multimedia and Expo (ICME)*, pages 1–6. IEEE, 2018.