

# A STRONG AND EFFICIENT BASELINE FOR VEHICLE RE-IDENTIFICATION USING DEEP TRIPLET EMBEDDING

Ratnesh Kumar<sup>1,\*</sup>, Edwin Weill<sup>1,\*</sup>, Farzin Aghdasi<sup>1</sup>, Parthasarathy Sriram<sup>1</sup>

<sup>1</sup>NVIDIA

\*E-mail: {ratneshk, eweill}@nvidia.com

Submitted: 30th September 2019; Accepted: 11th November 2019

## Abstract

In this paper we tackle the problem of vehicle re-identification in a camera network utilizing triplet embeddings. Re-identification is the problem of matching appearances of objects across different cameras. With the proliferation of surveillance cameras enabling smart and safer cities, there is an ever-increasing need to re-identify vehicles across cameras. Typical challenges arising in smart city scenarios include variations of viewpoints, illumination and self occlusions. Most successful approaches for re-identification involve (deep) learning an embedding space such that the vehicles of same identities are projected closer to one another, compared to the vehicles representing different identities. Popular loss functions for learning an embedding (space) include contrastive or triplet loss. In this paper we provide an extensive evaluation of triplet loss applied to vehicle re-identification and demonstrate that using the recently proposed sampling approaches for mining informative data points outperform most of the existing state-of-the-art approaches for vehicle re-identification. Compared to most existing state-of-the-art approaches, our approach is simpler and more straightforward for training utilizing only identity-level annotations, along with one of the smallest published embedding dimensions for efficient inference. Furthermore in this work we introduce a formal evaluation of a triplet sampling variant (batch sample) into the re-identification literature. In addition to the conference version [24], this submission adds extensive experiments on new released datasets, cross domain evaluations and ablation studies.

**Keywords:** convolutional neural networks, re-identification, triplet networks, siamese networks, embedding, hard data mining, contrastive loss

## 1 Introduction

Matching appearances of objects across multiple cameras is an important problem for many computer vision applications, *e.g.* object retrieval and object identification. This problem of object re-identification is closely related to object recognition and fine grained classification. In the realm of video understanding, most higher level algorithms such as event recognition and anomaly detection rely upon

*Multiple Camera Multiple Object Tracking (MC-MOT)*. An important component for a MC-MOT is an *object verification* (*i.e.* re-identification) module for expressing *confidence* to associate objects across multiple videos [40]. Re-identification approaches can also be used in a single camera setup, wherein the task would be to determine if the same object has re-appeared in the scene [23, 56, 48].

The task of vehicle re-identification is to identify the same vehicle across a camera network. With the deployment of camera sensors for traffic management and smart cities, there is an imminent need to perform vehicle search from video databases [38]. Previous works [47, 17] have shown that automatic recognition of license plates as a global unique identifier have given state-of-the-art identification performance. However in general traffic scenes at streets, license plates are practically invisible in many views to recognize due to their top view installations. Therefore, a vision-based re-identification has a great practical value in real world scenarios. Re-identification of objects is challenging due to significant appearance & viewpoint shifts, lighting variations and varied object poses. Figure 1 shows some typical challenging intra-class variations.



**Figure 1.** Each row is a separate identity (samples taken from VeRi [29] and VRIC [20] datasets).

Despite large intra-class variations for views, vehicle-model could be discerned from most views. Other challenges to this task include occlusion, motion blur (3rd row).

Compared to person and face re-identification, vehicle re-identification is a relatively understudied problem. A few of the unique characteristics pertaining to the problem of vehicle re-identification which make it a difficult task are:

- Multiple views of the same vehicle are visually diverse and semantically (i.e. color and model)

correlated, meaning that the same identity must be deduced no matter which viewpoint of the vehicle is given.

- In real world scenarios, a re-identification system is expected to extract subtle physical cues such as the presence of dust, written marks, or dents on vehicle surfaces, to be able to distinguish between vehicles which are the same color and model.
- The vehicle labels are less fine-grained than person (or face)-identity labels. Given that there are a finite number of vehicle colors and models, the diversity in a given dataset is less than that of a person or face re-identification dataset.

In order to match appearances of objects, firstly we need to obtain an embedding for the objects, also denoted as a feature vector or signature. A match is then performed by using a suitable distance metric expressing the closeness of two objects in an embedding space. A good embedding should be invariant to illumination, scale and viewpoint changes. Prior to the advancements in deep learning, most embedding learning approaches focus on handcrafting using mixture of multiple feature extractors and/or learning suitable ranking functions to minimize distance across objects of similar identities. Some of the notable approaches are [51, 3, 34, 33, 6, 26, 59].

*In this paper* we focus on the embedding part of the re-identification process and make the following **contributions**:

- Utilizing the recent advances in *sampling* informative data points for learning embedding for the person re-identification task [12], we extensively evaluate their application to the vehicle re-identification problem, and demonstrate state-of-the-art performance across diverse datasets on various performance metrics.
- We introduce a formal evaluation of a triplet sampling variant, *batch sample*, into the re-identification literature.
- In addition to the conference version [24] we add evaluations on three more recently proposed datasets with exhaustive *cross dataset* evaluation experiments establishing domain gap in these datasets.

The rest of the paper is organized as: in the following Section we provide an overview of related works and the subsequent Section will elaborate on triplet and contrastive losses, including popular sampling techniques to optimize these losses. Section 4 details on datasets and hyperparameters used for various experimental settings. Results and discussions are presented in Section 5.

## 2 Related Works

In recent years with the evolution of end-to-end learning using *Convolutional Neural Networks* (CNN), significant improvements have been made in feature representations using large amounts of training data. These approaches outperform all previous baselines using handcrafted features. A CNN learns hierarchical image features by stacking convolutional layers with downsampling layers. The outputs from one convolutional layer is fed to a non-linearity layer before being fed to the subsequent convolutional layer.

[4] proposed one of the first approaches to learn visual relationships using CNN. *Siamese CNN* [4] computes an embedding space such that similar examples have similar embeddings and vice versa. [5] uses *contrastive* loss on Siamese CNN to learn embedding for face verification. One of the recent prominent works using CNNs for learning face embedding [41] uses *triplet loss* to train a CNN for learning face embeddings for identification. While triplet loss considers three samples *jointly* for computing a loss measure, contrastive loss requires only two samples. Contrastive loss is computationally more efficient than triplet, however, several approaches [35, 9, 40, 2, 12, 14] have reported state-of-the-art performances using triplet loss. This superiority of triplet loss is attributed to the additional context using three samples. Section 3 in this paper elaborates on these losses.

Another method for obtaining an embedding for an object is utilizing a traditional softmax layer [57, 19], wherein a fully-connected (embedding) layer is added prior to the softmax-loss layer. Each identity is considered as a separate category and the number of categories is equal to the number of identities in the training set. Once the network is trained using classification loss (*e.g.* cross-entropy), the classification layer is stripped off and an embed-

ding is obtained from the new final layer of the network. [19] proposed a similar approach to learning vehicle embedding based on training a network for vehicle-model classification task. Since the network is not directly trained on embedding or metric learning loss, usually the performance of such a network is poor when compared to networks incorporating embedding loss. Cross entropy loss ensures separability of features but the features may not be discriminative enough for separating unseen identities. Furthermore learning becomes computationally prohibitive when considering datasets of *e.g.*  $10^6$  identities. Some recent works [56, 39, 42] unify classification loss with metric learning.

**Vehicle Classification:** Fine grained vehicle classification is a closely related problem to vehicle re-identification. Notable works for vehicle classification are [25, 8, 16, 27, 46, 32]. The general task is to predict vehicle *model*, *e.g.* BMW-i3-2016, Toyota-Camry-1996. Vehicle re-identification is a relatively finer grained problem than vehicle-model classification: a re-identification approach should be able to extract visual differences between two vehicles belonging to the same model category. The visual differences could include subtle cosmetic and color differences making this problem more difficult. Furthermore a re-identification method is expected to work without any *a priori* knowledge of all possible vehicle models in the city or a geographical entity.

**Vehicle Re-identification:** Some notable approaches prior to deep learning are [30, 59]. Popular deep learning approaches for vehicle re-identification are [52, 28, 58, 2, 29, 9, 43, 53, 61, 19, 62]. [29] proposed fusion of handcrafted features *e.g.* color, texture along with high level attribute feature obtained using CNN. [52] proposed a progressive refinement approach to searching query vehicles. A list of candidates is obtained for a query using embeddings from a siamese-CNN trained using contrastive loss. This list is then pruned using a siamese network to match license plates. In order to get reliable query for visually similar vehicles, authors factor in the usage of spatio-temporal distance comparison in addition to visual embedding distances.

[9] presents a structured deep learning loss comprising a classification loss term (based on vehicle model) as well as coarse and fine grained rank-

ing terms. [28] proposed a modification of triplet loss by replacing anchor samples with corresponding class center in order to suppress effects of using poor anchors. Furthermore the deep model is trained for both vehicle model classification and identity labels in a multi-level process. [58] focuses on the relationship between different vehicle images as multiple grains by using diverse vehicle attributes. The authors proposed ranking methods incorporated into multi-grain classification.

In a recent work [2], the authors propose to include group-based sub-clustering in a triplet loss framework. This helps in explicitly dealing with intra-class variations of vehicle identification problem. During training an online grouping method is used to cluster samples within each identity into disparate clusters. The authors demonstrate state-of-the-art results in different datasets. Moreover we also add recently proposed algorithms in evaluations.

[61] proposes to use a view-point synthesis approach to predict embedding for unknown views given a true view image. These synthetic embeddings for unknown views are generated using bi-directional LSTM [13]. The complete network is trained using a combination of contrastive, reconstruction and generative adversarial loss [7]. Similar to the objective of [61] for inferring a global feature vector using view-synthesis, authors in [62] propose a *viewpoint attentive multi-view* framework. Utilizing attentive [37] and adversarial loss, authors transform a single view feature into a global multi-view feature representation.

[53] develops a framework utilizing keypoint annotations on vehicles to learn viewpoint invariant features from a CNN. To further enhance the retrieval of matching vehicles the authors use probabilistic spatio-temporal regularization using random variables representing camera transition probabilities. The authors demonstrate superior results by adding this regularization during retrieval procedure. [43] formulate these *camera transition* probabilities by generating proposals of path (trajectories) and employing a LSTM and Siamese CNN to obtain a robust re-identification performance. FDA-net proposed by [31] introduced a novel hard data generator coupled with an end to end training with feature distance adversarial scheme.

In order to deal with lack of fine grained data to include vehicle attributes for re-identification, [49] proposed PAMTRI (Pose-Aware Multi-Task Re-Identification), which explicitly reason about vehicle pose and shape via keypoints, heatmaps and segments from pose estimation. Training is performed in multi-task learning fashion. [20] proposed a multi-scale matching approach (MSVR) to vehicle re-identification by learning more discriminative feature representations from multi-resolution images.

### 3 Loss functions for embedding

For a reliable re-identification of objects, the following are some desired characteristics of an embedding function:

- An embedding should be invariant to view-points, illumination and shape changes to the object.
- For a practical application deployment, computation of embedding and ranking should be efficient.

Consider a dataset  $\mathcal{X} = \{(x_i, y_i)\}_{i=1}^N$  of  $N$  training images  $x_i \in \mathbb{R}^D$  and their corresponding class labels  $y_i \in \{1 \cdots C\}$ . Re-identification approaches aim to learn an embedding  $f(x; \theta) : \mathbb{R}^D \rightarrow \mathbb{R}^F$  to map images in  $\mathbb{R}^D$  onto a feature (embedding) space in  $\mathbb{R}^F$  such that images of similar identity are metrically close in this feature space.  $\theta$  corresponds to the parameters of the learning function.

$$\theta^* = \underset{\theta}{\operatorname{argmin}} \mathcal{L}(f(\theta, \mathcal{X})). \quad (1)$$

Let  $D(x_i, x_j) : \mathbb{R}^F \times \mathbb{R}^F \rightarrow \mathbb{R}$  be a metric measuring distance of images  $x_i$  and  $x_j$  in embedding space. For simplicity we drop the input labels and denote  $D(x_i, x_j)$  as  $D_{ij}$ .  $y_{ij} = 1$  is both samples  $i$  and  $j$  belong to the same class and  $y_{ij} = 0$  indicates samples of different classes.

#### 3.1 Contrastive Loss

Contrastive loss (2) was employed in [5] for the face verification problem, wherein the objective is to verify if two presented faces belong to the same identity. This discriminative loss directly optimizes

(1) by encouraging all similar class distances to approach 0 while keeping all dis-similar class distances to be above a pre-defined threshold  $\alpha$ .

$$l_{contrastive}(i, j) = y_{ij}D_{ij}^2 + (1 - y_{ij})[\alpha - D_{ij}^2]_+. \quad (2)$$

Notice that the choice of  $\alpha$  is same for all dissimilar classes. This implies that for dissimilar identities, visually diverse classes are embedded in the same feature space as the visually similar ones. This assumption is stricter when compared to triplet loss, and restricts the structure of the embedding manifold thereby impairing discriminative learning. The training complexity is  $O(N^2)$  for a dataset of  $N$  samples.

### 3.2 Triplet Loss

Inspired from the seminal work on metric learning for nearest neighbor classification by [54], *facenet* [41] proposed a modification suited for retrieval tasks *i.e.* equation (3), termed: triplet loss. Triplet loss forces the data points from the same class to be closer to each other than a data point from any other class. Notice that contrary to contrastive loss in (2), triplet loss adds context to the loss function by considering both a positive and negative pair distances from the same point.

$$l_{triplet}(a, p, n) = [D_{ap} - D_{an} + \alpha]_+. \quad (3)$$

Training complexity of triplet loss is  $O(N^3)$  which is computationally prohibitive. High computational complexity of triplet and contrastive losses have motivated a host of sampling approaches for an efficient optimization.

### Dataset Sampling

As triplet and contrastive losses are computationally prohibitive for practical datasets, most proposed approaches resort to sampling *effective* data points for computing losses. This is important as computing loss over trivial data points could only impair convergence of the algorithm. In the context of vehicles, it will be more informative for a loss function to sample from different views (*e.g.* side or front-view) for the same identity, than considering samples of similar views repeatedly.

A popular sampling approach to find informative samples is *hard data mining*, and is employed

in many computer vision applications *e.g.* object detection. Hard data mining is a bootstrapping technique which is used in iterative training of a model, wherein at every iteration the current model is applied on a validation set to mine hard data on which this model is performing poorly. Only these hard data are then presented to the optimizer which increases the ability of the model to learn effectively and converge faster to an optimum. On the flip side, if a model is only presented with hard data, which could comprise outliers, its ability to discriminate outliers *w.r.t.* normal data would suffer.

In order to deal with the outliers during hard data sampling, *facenet* [41] proposed *semihard* sampling which mines moderate triplets that are neither too hard nor too trivial for getting meaningful gradients during training. [12] proposed an efficient and effective approach to mine samples directly on GPU. They emphasize on having good batch statistics to encourage mining efficient data samples. The authors construct a data batch by randomly sampling  $P$  identities from  $X$  and then randomly sampling  $K$  images for each identity, thus resulting in a batch size of  $PK$  images. In a batch size of  $PK$  images, the authors [12] proposed two sampling techniques, namely **batch hard** (BH) (also in [36]) and **batch all** (BA). Another sampling technique **batch sample** (BS) is actively discussed in the implementation webpage of [12], however to the best of our knowledge we could not find a formalized study and evaluation for this sampling technique.

[40] unifies different batch sampling techniques in [12] under one expression. Let  $a$  be an anchor sample and  $N(a)$  and  $P(a)$  represent a subset of negative and positive samples for the corresponding anchor  $a$ . The triplet loss can then be written as:

$$l_{triplet}(a) = [\alpha + \sum_{p \in P(a)} w_p D_{ap} - \sum_{n \in N(a)} w_n D_{an}]_+. \quad (4)$$

With respect to an anchor sample  $a$ :  $w_p$  represents the weight (importance) of positive sample  $p$  and similarly  $w_n$  signifies the importance of the negative sample  $n$ .

The total loss in an epoch is then obtained by

$$\mathcal{L}(\theta; X) = \sum_{all \ batches} \sum_{a \in B} l_{triplet}(a).$$


**Table 1.** Various ways of mining good samples in a batch, for better optimization of embedding loss

Sampling variant	Positive weight: $w_p$	Negative weight: $w_n$	Comments
Batch all (BA)	1	1	Uniformly weighted
Batch hard (BH)	$[x_p == \arg \max_{x \in P(a)} D_{ax}]$	$[x_n == \arg \min_{x \in N(a)} D_{ax}]$	Hardest sample
Batch sample (BS)	$[x_p == \text{multinomial} \{D_{ax}\}]$	$[x_n == \text{multinomial} \{-D_{ax}\}]$	Multinomial sampling
Batch weighted (BW)	$\frac{e^{D_{ap}}}{\sum_{x \in P(a)} e^{D_{ax}}}$	$\frac{e^{-D_{an}}}{\sum_{x \in N(a)} e^{-D_{ax}}}$	Adaptive weights

Table 1 summarizes different ways of sampling positives and negatives. We formalize BS method in this regime. BH is hard data mining in the batch, using only the hardest positive and negative samples for every anchor. BA is a straightforward sampling which gives uniform weights to all samples. Uniform weight distribution can ignore the contribution of important tough samples as these samples are typically outnumbered by the trivial easy samples. In order to mitigate this issue with BA, [40] employs a weighting scheme **batch weighted** (BW), wherein a sample is weighted based on its distance from the corresponding anchor, thereby giving more importance to the informative harder samples than trivial samples.

BS uses the distribution of anchor-to-sample distances to mine a positive and negative data for an anchor. This technique thereby avoids sampling outliers when compared with BH, and also hopes to find out the most relevant sample as the sampling is done using distances-to-anchor distribution.

A sample illustration of our approach to compute an embedding is shown in Figure 2.

**Figure 2.** During inference an image is passed through a CNN and the final layer provides the corresponding embedding.

In the following Sections, we evaluate the embedding losses, along with the sampling variants presented in Table 1.

## 4 Experiments

For our evaluation purposes we use three popular publicly available datasets: VeRi, VehicleID and PKU-VD.

**VeRi:** This dataset is proposed by [29] and is one of the main datasets used in vehicle re-identification literature for comparative study. This dataset encompasses 40,000 bounding box annotations of 776 cars (identities) across 20 cameras in traffic surveillance scenes. Each vehicle is captured in 2-18 cameras in various viewpoints and varying illuminations. Notably the viewpoints are not restricted to only front/rear but also side views, thereby making it one of the challenging datasets. The annotations include make and model of vehicles, color and inter-camera relations and trajectory information.

**VehicleID:** This dataset [28] comprises 221,763 bounding boxes of 26,267 identities, captured across various surveillance cameras in a city. Annotations include 250 vehicle models and this dataset has an order of magnitude more images than VeRi dataset. However the viewpoints only include front and rear views for vehicles.

**PKU-VD:** [58] proposed a large dataset for fine grained vehicle analysis including re-identification and classification. To this date this is the largest dataset comprising about *two million* images and their fine grained labels including vehicle model and color. This dataset is split into two sub-datasets, namely **VD1** and **VD2** based on cities from which they were captured. The images in VD1 are captured from higher resolution cameras than VD2. There are about 71k and 36k identities in VD1 and VD2, respectively.

**CityFlow:** Presented at CVPR 2019 [50], this large scale dataset is the biggest city-wide dataset collected in USA. The dataset comprises annotations for the tasks of re-identification, multi-camera-

multi-target vehicle tracking, containing more than 200K annotated bounding boxes covering a wide range of scenes, viewing angles, vehicle models, and urban traffic flow conditions. For our purposes we utilize the subset of the dataset annotated for image based re-identification. This set contains 333 identities over 36.935 bounding boxes for training, and 333 identities for testing.

**Veri-Wild:** This dataset [31] was also presented at CVPR 2019, comprising about 30,000 identities over 277,797 images for training. The test set is split into three subsets: small, medium and large. This large city-wide dataset is captured from a network of 174 cameras spanning 200 square kilometers.

**VRIC:** This dataset [20] Vehicle Re-identification in Context, was presented at GCPR 2018. VRIC is uniquely characterised by vehicle images subject to more realistic and unconstrained variations in resolution (scale), motion blur, illumination, occlusion, and viewpoint. It contains 60,430 images of 5,622 vehicle identities captured by 60 different cameras in both day-time and night-time.

#### 4.1 Training and Hyperparameters

For all our experiments, we fix our backbone or meta-architecture to *mobilenet-v1* [15] owing to its better efficiency (parameters, speed) as compared to ResNet-variants [11] and VGG [45]. The imagenet [18] retrieval accuracy for these architectures are in similar ranges. Mobilenet-v1 has 569 million Multiply Accumulates (MACs) which measures the number of fused multiplication and addition operations. This architecture has 4.24 million parameters and achieves a top-1 accuracy of 70.9 on imagenet’s image classification benchmark, with input image size of 224x224.

We use Adam optimizer [22] with default hyperparameters ( $\epsilon = 10^{-3}$ ,  $\beta_1 = 0.9$ ,  $\beta_2 = 0.999$ ). Depending upon if the training is done from scratch or fine-tuned using an imagenet [18] based trained model, we employ different learning rate schedulers. When training from scratch, we use standard learning rate of 0.001. We reduce this rate to 0.0003 when using an imagenet based pre-trained model. For online data augmentation a standard image-flip operation is used. We use Nvidia’s Volta GPU for hardware and Tensorflow [1] as the software platform.

We replace the margin  $\alpha$  in triplet loss (4) by *softplus* function:  $\ln(1 + \exp(\cdot))$  which avoids the need of tuning this margin [12]. For contrastive loss we follow standard practice of hard margin of 1.0. Using a softplus function produced poorer results for contrastive loss.

For the *batch construction*, unless otherwise specified, we follow the default batch sizes as in [12, 40]. A batch consists of 18 (P) randomly chosen identities, and for each identity, 4 (K) samples are chosen randomly, thereby selecting a total of 72 (PK) images. Samples are chosen such that we iterate over all train set during the course of an epoch. Following the standards in face-verification and person re-identification [40], [41] we set the embedding dimension to *128 units*.

#### 4.2 Evaluation Metrics

We use mean-average-precision (*mAP*) and *top-k* accuracy for evaluating and comparing our presented approaches. In a typical re-identification evaluation setup, we have a query set and a gallery set. For each vehicle in a query set the aim is to retrieve a similar identity from the test set (*i.e.* gallery set). *AP*(*q*) for a query image *q* is defined as

$$AP(q) = \frac{\sum_k P(k) \times \delta_k}{N_{gt}(q)},$$

where *P*(*k*) represents precision at rank *k*, *N<sub>gt</sub>*(*q*) is the total number of true retrievals for *q*.  $\delta_k$  is 1 when the matching of query image *q* to a test image is correct at rank  $\leq k$ . *mAP* is then computed as average over all query images

$$mAP = \frac{\sum_q AP(q)}{Q},$$

where *Q* is the total number of query images.

### 5 Results and Discussions

We present our results on the datasets mentioned in the previous Section. Different datasets have different ways of constructing test sets which we elaborate in the respective Sections. Each model presented below is trained separately on the corresponding dataset using its standard train set.

## 5.1 VeRi

We follow the standard evaluation protocol by [52]. The total number of query images is 1,678 while the gallery set comprises 11,579 images. For every query image, the gallery set contains images of same query-identity but taken from *different* cameras. This is an important evaluation exclusion as in many cases the same camera samples would contain visually similar samples for the same vehicle.

**Table 2.** VeRi accuracy results (%) using triplet and contrastive loss for different batch sampling variants outlined in Table 1.

Sampling	mAP	top-1	top-2	top-5
<b>Triplet, Not-Normalized</b>				
BH	65.10	87.25	91.54	94.76
BA	66.91	90.11	<b>93.38</b>	96.01
BS	<b>67.55</b>	<b>90.23</b>	92.91	96.42
BW	67.02	89.99	93.15	<b>96.54</b>
<b>Triplet, Normalized</b>				
BH	53.72	72.65	80.27	86.83
BA	27.60	42.91	53.16	67.76
BS	33.79	48.75	58.64	73.54
BW	44.29	60.91	69.85	80.63
<b>Contrastive, Normalized</b>				
BH	59.21	80.51	85.52	90.64
BS	52.09	71.51	78.84	86.95
<b>Contrastive, Not-Normalized</b>				
BH	56.84	75.33	82.30	90.29
BS	48.85	65.49	74.55	85.76

Table 2 summarizes our results for various sampling configurations, and we can draw following inferences:

- Adding a normalized layer performs poorly for the triplet loss. This is also reported by [12] wherein using a normalized layer could result in collapsed embeddings.
- Siamese (contrastive) loss under performs relative to triplet loss. We attribute this to the additional context provided by using both positive and negative samples in the same term for the triplet loss [35].
- For the best performing set, *i.e.* triplet loss with no-normalization layer: all four sampling variants reach about similar accuracy ranges, with BS outperforming others in a close range.

- Figure 3 shows some visual results with embeddings learned from batch-sampling triplet loss. Good top-k retrievals indicate stability of our embeddings across different views and cameras. Notice that query and gallery images are constrained to be from different cameras following the standard evaluation protocol.

### Comparison to the state-of-the-art approaches:

Table 3 outlines comparisons with the state-of-the-art approaches. Notice that our approach performs close to the state-of-the-art. GSTE [2] achieves better top-k accuracy but in terms of mAP our approach performs better indicating robustness at all ranks. Furthermore GSTE [2] has an embedding dimension of 8x more (*i.e.* 1024) than ours, and GSTE includes a complicated training process which requires tuning an additional intra-class clustering parameter. PAMTRI (All) achieves the best results, while we outperform its basic version PAMTRI (RS). PAMTRI (RS) uses mix of real and synthetic data for learning embedding, while PAMTRI (All) additionally utilizes vehicle keypoints and attributes in a multi-task learning framework.

**Table 3.** Comparison of various proposed approaches on VeRi dataset. (\*) indicates the usage of spatio-temporal information.

Method	mAP	top-1	top-5
BS (Ours)	67.55	90.23	96.42
GSTE [2]	59.47	<b>96.24</b>	<b>98.97</b>
VAMI [62]	50.13	77.03	90.82
VAMI+ST * [62]	61.32	85.92	91.84
OIFE [53]	48.00	89.43	-
OIFE+ST *[53]	51.42	92.35	-
PROVID * [52]	27.77	61.44	78.78
Path-LSTM * [43]	58.27	83.49	90.04
PAMTRI (RS) [50]	63.76	90.70	94.40
PAMTRI (All) [49]	<b>71.88</b>	92.86	96.97
FDA-Net [31]	55.49	84.27	92.43
MSVR [20]	49.30	88.56	-
AAVER [21]	61.18	88.97	94.70

AAVER [21] is a recent work presented at ICCV 2019. Authors construct a dual path network for extracting global and local features. These are then concatenated to form a final embedding. The proposed embedding loss is minimized using identity and keypoint orientation annotations. For a fair



comparison, we include the results from AAVER without re-ranking. Our straightforward approach outperforms AAVER, and noticeably the best results by AAVER uses Resnet-101 with embedding dimension of 2048.

The VeRi dataset includes spatio-temporal (ST) information and [62, 61, 53, 52] utilize ST information in either embedding or in retrieval stages. Noticeably without using any ST information, we outperform these approaches using ST. Contrary to us, OIFE [53] requires extra annotations of keypoints during training for their orientation invariant embedding learning. Training procedure for VAMI [62] include *generative adversarial network* (GAN) and multi-view attention learning. Path-LSTM [43] employ generation of several path-proposals for their spatio-temporal regularization and requires an additional LSTM to rank these proposals. It is worth noting that our training procedure is more straightforward than most of the approaches presented in Table 3, with an efficient embedding dimension of 128. Table 4 outlines some important differences *w.r.t.* competitive approaches.

**Table 4.** Summary of some important hyperparameters and labeling used during training. **ED** indicates embedding dimension. **K**, **A** denotes keypoints and attributes respectively. OIFE merges four datasets to form one large training set. Notice that our ED is the least among other approaches.

PAMTRI’s both approaches use synthetic data (details in Section 5.4).

Method	ED	Annotations
Ours	128	ID
GSTE [2]	1024	ID
VAMI [62]	2048	ID + A
OIFE [53]	256	ID + K
MGR [58]	1024	ID + A
ATT [58]	1024	ID + A
C2F [9]	1024	ID + A
CLVR [19]	1024	A
PAMTRI (All)* [49]	1024	ID + K + A
MSVR [20]	2048	ID
FDA-Net [31]	1024	ID
AAVER [21]	2048	ID + K

Referring to the best results in Table 2, in the subsequent Sections we consider only triplet loss without embedding-normalization.

## 5.2 PKU-VD

PKU-VD is a large dataset combining two sub-datasets, VD1 and VD2. Both of these comprise about 400k training images. The test set of each of the sub-dataset is split into *three* reference sets: small, medium and large. Table 5 shows the number of test images in each sub-dataset. For evaluation, we use the same dataset files for each reference set as provided by the authors [58] of this dataset.

**Table 5.** Number of images in each reference test-set.

Dataset	Small	Medium	Large
VD1	106,887	604,432	1,097,649
VD2	105,550	457,910	807,260

**Table 6.** mAP (%) for retrievals on various reference sets. Training is performed from scratch *without* using pretrained weights.

Method	Small	Medium	Large
<b>VD1</b>			
BW (18x16)	<b>87.48</b>	<b>67.28</b>	<b>58.77</b>
MGR [58]	79.10	58.30	51.10
<b>VD2</b>			
BW (18x16)	<b>84.55</b>	<b>69.87</b>	<b>63.64</b>
MGR [58]	74.70	60.60	55.30

**Table 7.** mAP (%) for retrievals on various reference sets of different sizes. Training is performed without pretrained weights with batch size of 18x16.

Dataset, Sampling	Small	Medium	Large
<b>No pretrained weights</b>			
VD1, BA	85.02	62.84	54.68
VD1, BH	0.00	0.00	0.00
VD1, BS	87.24	66.62	58.26
VD2, BA	83.39	68.58	62.34
VD2, BH	0.00	0.00	0.00
VD2, BS	83.30	68.45	62.36

Compared to VeRi and VehicleID datasets, PKU-VD dataset has an order of magnitude more images, hence a deep network can be trained from scratch on this dataset. Furthermore with more intra-class samples, one can increase the batch size



**Figure 3.** Qualitative results on VeRi dataset using *BS* based triplet embedding. Each row indicates query image and top-10 retrievals for this query image. Red border indicates an incorrect retrieval and Green indicates a correct retrieval. These demonstrate good embedding quality as the top retrievals include different views and cameras.

of triplets. Tables 6, 7 and 8 show results for various configurations. For the BW sampling in Table 6, the numerics following illustrate the  $P$  and  $K$  values, described previously, which create the batch. Table 7 adds results for the other three sampling variants when training from scratch. Table 8 shows results with default setting (batch size of  $18 \times 4$  and a pretrained mobilenet-v1).

**Table 8.** mAP (%) for retrievals on various reference sets of different sizes. Training is performed using imagenet pretrained weights with default batch size of  $18 \times 4$ .

Dataset, Sampling	Small	Medium	Large
<b>With pretrained weights</b>			
VD1, BW	<b>82.66</b>	60.15	52.10
VD1, BS	81.36	58.91	50.68
VD1, BA	79.46	56.79	49.26
VD1, BH	82.04	<b>60.40</b>	<b>52.17</b>
VD2, BW	<b>80.93</b>	<b>65.44</b>	<b>58.94</b>
VD2, BS	75.52	58.35	51.71
VD2, BA	70.07	50.56	43.46
VD2, BH	78.95	62.32	55.86

Using more triplets in the batch improves the accuracy, which is intuitively satisfying. Noticeably using the hardest sample (BH) does not kick-off the training (*c.f.* Table 7). This is expected and also noted in [41], as with BH due to random initialization, the network never learns any understanding to separate hard data from easy samples. One way to deal with this is to start training with a few identities in a multi-class setting in-order to pre-train the network and then proceed with the standard BH procedure. Alternatively one could start from an imagenet trained network (*c.f.* Table 8). Interestingly training from scratch results in a better accuracy than training from a pre-trained network. This observation is similar to [10], wherein the authors report using pre-trained weights (from imagenet) obtains similar results on large datasets when random initialization is used (for the base CNN architecture).

BW sampling with batch size of  $18 \times 16$  outperforms the precious state-of-the-art by [58]. Multi-grain ranking (MGR) uses permutation probability based ranking method and include vehicle attributes during training process. Noticeably our training procedure is straightforward without using vehicle attributes. Furthermore MGR uses an embedding

dimension of 1024 as opposed to 128 for our embedding, thus calling for higher computation cost during inference in [58].

### 5.3 VehicleID

VehicleID [28] is a larger dataset than VeRi containing front and rear views for the vehicles. We follow the standard evaluation protocol of [28] and provide results on four reference *query* sets. Reference sets: small, medium, large and X-large contain 800, 1600, 2400 and 13164 identities, respectively. For each reference set, an exemplar for an identity is randomly chosen, and a *gallery* set is constructed. This process is repeated ten times to obtain *averaged* evaluation metrics. For training we use mobilenet network, pretrained using imagenet dataset, without normalization-layer for embedding. Similarly to the PKU-VD dataset training we set the batch size ( $PK$ ) to  $18 \times 16$  images. For the sake of completeness we provide the results with default PK batch size of  $18 \times 4$ .

**Table 9.** Accuracy results on VehicleID using mAP metric (%). Batch size for our experiments is set to  $18 \times 16$  samples.

Method	Small	Medium	Large	X-Large
BA	84.65	79.85	75.95	59.74
BS	<b>86.19</b>	<b>81.69</b>	<b>78.16</b>	<b>62.41</b>
BW	85.92	81.41	78.13	62.12
BH	85.59	80.76	76.87	60.33
C2F [9]	63.50	60.00	53.00	-
GSTE [2]	75.40	74.30	72.40	-
ATT [58]	62.80	62.30	58.60	-
CCL [28]	54.60	48.10	45.50	-
FDA-Net [31]	-	65.33	61.84	-

**Table 10.** Accuracy results on VehicleID using mAP metric (%). This is with default PK batch size of ( $18 \times 4$ ).

Method	Small	Medium	Large	X-Large
BA	81.90	76.57	72.60	54.95
BS	84.17	79.05	75.52	59.10
BW	<b>84.90</b>	<b>80.80</b>	<b>77.20</b>	<b>60.92</b>
BH	83.34	78.72	75.02	57.97

Tables 9, 10 and 11 show comparative results for mAP and top-k metrics, respectively. Similarly to the PKU-VD results, using a larger batch size increases the retrieval rankings, however the margin of improvement is smaller. This could be due to

limited variability in this dataset in terms of viewpoints and number of vehicle-models, owing to which increasing the batch size does not necessarily increase informative statistics.

**Table 11.** Results on VehicleID dataset using top-k metric (%). Batch size for our experiments is set to 18x16 samples.

Method	Small	Medium	Large	X-Large
<b>Top-1</b>				
BA	76.69	71.20	66.71	50.22
BS	<b>78.80</b>	73.41	69.33	<b>53.07</b>
BW	78.49	73.10	69.41	52.82
BH	77.90	72.14	67.56	50.67
OIFE [53]	-	-	67.00	-
OIFE+ [53]	-	-	68.00	-
VAMI [62]	63.12	52.87	47.34	-
CCL [28]	49.00	42.80	38.20	-
C2F [9]	61.10	56.20	51.40	-
GSTE [2]	75.90	<b>74.80</b>	<b>74.00</b>	-
CLVR [19]	62.00	56.10	50.60	-
MSVR [20]	-	-	63.02	-
FDA-Net [31]	-	59.84	55.53	-
AAVER [21]	74.69	68.62	63.54	-
<b>Top-5</b>				
BA	95.26	91.17	87.75	70.48
BS	<b>96.17</b>	<b>92.57</b>	<b>89.45</b>	<b>73.06</b>
BW	95.83	92.48	89.36	72.72
BH	95.74	92.03	88.81	71.23
OIFE [53]	-	-	82.90	-
VAMI [62]	83.25	75.12	70.29	-
CCL [28]	73.50	66.80	61.60	-
C2F [9]	81.70	76.20	72.20	-
GSTE [2]	84.20	83.60	82.70	-
CLVR [19]	76.00	71.80	68.00	-
MSVR [20]	-	-	73.05	-
FDA-Net [31]	-	77.09	74.65	-
AAVER [21]	93.82	89.95	85.64	-

BS and BW outperform other sampling variants, including all state-of-the-art approaches in the mAP metric. Table 4 and Section 5.1 summarizes important differences of state-of-the-art approaches *w.r.t.* our approach. GSTE [2] achieves better performance in terms of top-1 accuracy, but their accuracy drops for top-5. Lower mAP and top-5 indicates GSTE’s sub-par retrieval performances for ranks  $k > 1$ . OIFE+ [53] achieves close accuracy in top-5 to ours. As opposed to our approach, OIFE+ requires keypoint annotations and a separate metric learning module from [60]. Furthermore OIFE combines VeRi, VehicleID, CompCars [32] and Cars21k [46] into one large train set. Contrary

to our method, other approaches [9, 28, 58], all utilize model annotations (in addition to identity annotations) from the training set for re-identification. Similarly to OIFE [53], AAVER [21] uses identity and orientation information for training the embedding network.

## 5.4 CityFlow dataset

**Table 12.** Comparison of proposed and state-of-the-art methods on CityFlow dataset.

Method	mAP	top-1	top-5
BA	31.72	49.62	65.02
BH	32.44	48.38	65.25
BS	31.73	49.05	63.12
BW	31.25	50.10	64.92
PAMTRI (All) [49]	<b>40.39</b>	<b>59.70</b>	<b>70.91</b>
PAMTRI (RS) [49]	31.41	50.37	61.48

We consider the image-to-image search (or image based re-identification) subset of this dataset. The query set consists of 1052 images and the test set comprises 333 identities over 18,290 images. Table 12 shows results of various sampling approaches and comparisons with state-of-the-art approaches. Our straightforward approach to training outperforms PAMTRI (RS) [49] uses synthetic data in conjunction with real data. Figure 4 shows qualitative assessment on our best performing sampling variant on this dataset. PAMTRI (All) performs best by a significant margin, utilizes annotations for keypoints and vehicle-attributes in a multi-task learning employing additional synthetic data. Furthermore the embedding dimension of PAMTRI is 1024 as opposed to 128 for our approaches.

## 5.5 VRIC

**Table 13.** Comparison of various approaches on VRIC dataset. We use our default network training scheme.

Method	mAP	top-1	top-5
BA	75.11	64.18	89.4
BH	77.99	67.77	91.32
BS	76.78	66.83	<b>90.64</b>
BW	<b>78.55</b>	<b>69.09</b>	90.54
MSVR [20]	-	46.61	65.58



**Figure 4.** Qualitative results on CityFlow dataset. Red border indicates an incorrect retrieval and Green indicates a correct retrieval. Last row shows a failure case - due to a subtle shape mismatch on the trunk, and the presence of background-structure leading to poor embedding of the foreground object. In future work we would investigate incorporating segmentation information in training and retrievals.

[20] proposed this dataset comprising large variations in scale, motion, illumination, occlusion and viewpoint. Authors [20] also proposed a multi-scale re-identification algorithm, termed MSVR.

Following the standard in [20], the query and probe splits are set to 2811 identities. Evaluation results from various sampling techniques are presented in Table 13. Our standard experimental settings outperform MSVR by a good margin, while also being efficient *w.r.t.* to the embedding dimension (128 vs. 2048 for MSVR).

## 5.6 Veri-Wild

**Table 14.** Results on Veri-Wild dataset using **top-k** metric (%) and **mAP**. Batch size for our experiments is set to 18x4 samples. Default experimental settings were used (embedding dimension at 128).

Method	Small	Medium	Large
<b>Top-1</b>			
BA	82.83	78.06	69.72
BH	83.30	76.90	69.10
BS	82.90	77.68	69.59
BW	<b>84.17</b>	78.22	69.99
FDA-Net [31]	64.03	57.82	49.43
<b>Top-5</b>			
BA	95.27	93.02	88.32
BH	95.20	92.66	87.74
BS	95.00	92.90	87.89
BW	<b>95.30</b>	93.06	88.45
FDA-Net [31]	82.80	78.34	70.48
<b>mAP</b>			
BA	68.21	60.69	49.28
BH	69.37	61.47	50.27
BS	68.79	61.11	49.79
BW	<b>70.54</b>	62.83	51.63
FDA-Net [31]	35.11	29.80	28.78

Veri-Wild [31] is the largest dataset as of CVPR 2019. This dataset comprises 416,314 vehicle images of 40,671 identities. Evaluation on this dataset is split across three subsets: small, medium and large; comprising 3000, 5000 and 10,000 identities respectively (in probe and gallery sets).

In addition to this large scale dataset, [31] also proposes a novel adversary algorithm termed FDA-Net, which generates online hard-negatives to guide re-identification training. FDA-Net uses an embed-

ding dimension of 1024 and is based on VGG [44] CNN meta-architecture.

Table 14 demonstrates our results on this dataset. Compared to FDA-Net [31], we achieve significantly better results on all test sets. Qualitative visual results are shown in Figure 5.

## 6 Cross domain evaluation

Owing to the importance of vehicle re-identification task, there is a host of dataset proposed in the recent years. The presented approach in this paper achieves state-of-the-art results on most of these datasets. However we train a separate network in a fully supervised fashion on each dataset, which is cumbersome and hinders wider deploy-ability. This can be improved by approaches such as domain adaptation or mixing domains to make a wider domain which can be used as a baseline, hopefully performing reasonably well on a large set of test-domains. In order to design and develop a model which is widely applicable, it is imperative to understand the domain gap between the datasets. To this end we do a full extent cross domain evaluation of the models developed for each dataset. The results are presented in Table 6. Several inferences can be drawn from this cross domain evaluation:

- Unsurprisingly, accuracy is at best when a model is trained and tested on the same dataset.
- Datasets: VD1, VD2, VehicleID demonstrates less domain gap between them (as opposed to comparing with other datasets).
- VeRi and VehicleID datasets seem to perform reasonably well when tested on each other. Similarly Veri-Wild has reasonable performance when tested on VehicleID and VeRi.
- CityFlow and VRIC datasets appear most challenging in terms of cross domain testing. This gap could be due to the fact that these two datasets were captured in USA (different cities), as opposed to other datasets which were collected outside USA. VRIC is a subset of UADE-TRAC dataset [55].



**Figure 5.** Qualitative results on VeRi-Wild dataset. Each row indicates query image and top-10 retrievals for this query image. Red border indicates an incorrect retrieval and Green indicates a correct retrieval. We observe that embeddings are robust to viewpoints. Most incorrect results look visually correct, and underlines the need for incorporating temporal aspect in retrievals.

**Table 15.** Cross domain evaluations: **S**, **T** indicates *source* and *target* dataset domains respectively. Best **mAP** numbers are in red and the second best accuracy is colored in blue. For Veri-Wild and VehicleID, we use the *large* test set for cross domain testing.

S (↓) T (→)	VeRi		VehicleID		VD-1		VD-2		CityFlow		VRIC		Veri-Wild	
	mAP	Top-1	mAP	Top-1	mAP	Top-1	mAP	Top-1	mAP	Top-1	mAP	Top-1	mAP	Top-1
CityFlow	18.29	48.81	27.49	20.75	7.70	14.15	6.09	15.15	<b>32.80</b>	51.05	8.61	4.70	8.05	30.92
VD-1	16.47	59.30	<b>64.70</b>	60.99	<b>58.77</b>	57.40	<b>52.53</b>	62.80	2.05	10.08	2.36	0.92	<b>25.16</b>	58.96
VD-2	13.31	52.21	64.11	60.03	<b>41.21</b>	48.75	<b>63.64</b>	69.55	1.70	8.75	3.09	1.46	18.33	50.11
VehicleID	24.75	57.75	<b>78.41</b>	69.63	27.32	35.35	39.49	49.35	2.60	10.36	3.75	1.85	18.81	48.76
VeRi	<b>67.55</b>	90.23	40.04	30.69	10.60	18.25	8.94	18.15	7.41	14.16	<b>16.67</b>	10.17	13.17	41.5
VRIC	7.37	20.08	15.76	11.67	4.89	10.40	3.85	10.60	1.56	3.80	<b>78.55</b>	69.09	3.2	16.56
Veri-Wild	<b>46.93</b>	74.49	63.42	51.01	17.17	23.45	16.90	25.05	<b>6.80</b>	20.44	7.53	4.52	<b>51.63</b>	69.99

## 7 Conclusion and Future Work

In this paper we propose a strong baseline for vehicle re-identification using the best practices in learning triplet embedding [12]. The core ideas behind this set of best practices lie in constructing a batch to facilitate extracting meaningful statistics in order to guide training and convergence. We introduced a formal exposition and evaluation of a sampling variant *batch sample* to the re-identification literature.

We compared our baselines with the state-of-the-art approaches on three datasets and outperform almost all of them in a wide range of evaluation criteria. The sampling variants: *batch sample* and *batch weighted* proved generally more effective and robust than *batch hard* and *batch all*.

In terms of parameters’ for training, the default setting (i.e. batch size of 18x4 samples, mobilenet-v1 with imagenet weights) performs reasonably well and outperforms most competing approaches.

We hinged our research on the belief that despite the intra-class variations, the identity of a vehicle is less fine grained than other object re-identification task, *e.g.* person re-identification. Our results demonstrate this by using the recent advances in embedding learning, we can push the frontiers of vehicle re-identification much further without using any spatio-temporal information. On the other hand, two vehicles of exactly the same color and model (with subtle or no discerning marks, *e.g.* last row in Figure 3) would be very difficult to distinguish without any spatio-temporal information. Incorporating spatio-temporal information along with other attributes in an effective

manner is an important contribution as future work. Datasets such as CityFlow, VERI-Wild, VeRi and AI-CITY challenges [38] aim to further enhance the state-of-the-art by providing spatio-temporal annotations.

Our approach is amenable to real time deployment across camera networks. During inference, Mobilenet-v1 takes about 0.60 ms/image and 0.53 ms/image on batch sizes of 8 and 64 respectively, on Nvidia Titan Xp, using floating point precision (FLOAT32). The performance throughput can be further optimized on Nvidia GPU(s) using TensorRT<sup>1</sup> and/or choosing other precision formats *e.g.* INT8.

## References

- [1] M. Abadi, A. Agarwal, P. Barham, E. Brevdo, Z. Chen, C. Citro, G. S. Corrado, A. Davis, J. Dean, M. Devin, S. Ghemawat, I. Goodfellow, A. Harp, G. Irving, M. Isard, Y. Jia, R. Jozefowicz, L. Kaiser, M. Kudlur, J. Levenberg, D. Mané, R. Monga, S. Moore, D. Murray, C. Olah, M. Schuster, J. Shlens, B. Steiner, I. Sutskever, K. Talwar, P. Tucker, V. Vanhoucke, V. Vasudevan, F. Viégas, O. Vinyals, P. Warden, M. Wattenberg, M. Wicke, Y. Yu, and X. Zheng. TensorFlow: Large-scale machine learning on heterogeneous systems, 2015. Software available from tensorflow.org.
- [2] Y. Bai, Y. Lou, F. Gao, S. Wang, Y. Wu, and L. Duan. Group Sensitive Triplet Embedding for Vehicle Re-identification. IEEE Transactions on Multimedia, 2018.
- [3] S. Bak, M. S. Biagio, R. Kumar, V. Murino, and F. Bremond. Exploiting Feature Correlations by Brownian Statistics for People Detection and

<sup>1</sup><https://developer.nvidia.com/tensorrt>



- Recognition. *IEEE Transactions on Systems, Man, and Cybernetics*, 2017.
- [4] J. Bromley, J. W. Bentz, L. Bottou, I. Guyon, Y. Lecun, C. Moore, E. Säcker, and R. Shah. Signature Verification Using a “Siamese” Time Delay Neural Network. *International Journal of Pattern Recognition and Artificial Intelligence* 1993.
- [5] S. Chopra, R. Hadsell, and Y. LeCun. Learning a similarity metric discriminatively, with application to face verification. In *CVPR*, 2005.
- [6] M. Farenzena, L. Bazzani, A. Perina, V. Murino, and M. Cristani. Person Re-Identification by Symmetry-Driven Accumulation of Local Features. In *CVPR*, 2010.
- [7] I. J. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, and Y. Bengio. Generative Adversarial Networks. In *NIPS*, 2014.
- [8] H. Z. Gu and S. Y. Lee. Car model recognition by utilizing symmetric property to overcome severe pose variation. *Machine Vision and Applications*, 2013.
- [9] H. Guo, C. Zhao, Z. Liu, J. Wang, and H. Lu. Learning Coarse-to-Fine Structured Feature Embedding for Vehicle Re-Identification. In *AAAI*, 2018.
- [10] K. He, R. B. Girshick, and P. Dollár. Rethinking imagenet pre-training. *CoRR*’18
- [11] K. He, X. Zhang, S. Ren, and J. Sun. Deep Residual Learning for Image Recognition. In *CVPR*, 2016.
- [12] A. Hermans, L. Beyer, and B. Leibe. In Defense of the Triplet Loss for Person Re-Identification. In *CoRR*, 2017.
- [13] S. Hochreiter and J. Schmidhuber. Long Short-Term Memory. *Neural Computation*, 1997.
- [14] E. Hoffer and N. Ailon. Deep metric learning using triplet network. In *ICLR Workshops*, 2015.
- [15] A. G. Howard, M. Zhu, B. Chen, D. Kalenichenko, W. Wang, T. Weyand, M. Andreetto, and H. Adam. MobileNets: Efficient Convolutional Neural Networks for Mobile Vision Applications. In *CVPR*, 2017.
- [16] Q. Hu, H. Wang, T. Li, and C. Shen. Deep CNNs with Spatially Weighted Pooling for Fine-Grained Car Recognition. *IEEE Transactions on Intelligent Transportation Systems*, 2017.
- [17] V. Jain, Z. Sasindran, A. Rajagopal, S. Biswas, H. S. Bharadwaj, and K. R. Ramakrishnan. Deep automatic license plate recognition system. *ICVGIP*, 2016.
- [18] Jia Deng, Wei Dong, R. Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. ImageNet: A large-scale hierarchical image database. In *CVPR*, 2009.
- [19] A. Kanaci, X. Zhu, and S. Gong. Vehicle Re-Identification by Fine-Grained Cross-Level Deep Learning. In *BMVC*, 2017.
- [20] A. Kanaci, X. Zhu, and S. Gong. Vehicle re-identification in context. In *Pattern Recognition - 40th German Conference, GCPR 2018, Stuttgart, Germany, September 10-12, 2018, Proceedings*, 2018.
- [21] P. Khorramshahi, A. Kumar, N. Peri, S. S. Rambhatla, J.-C. Chen, and R. Chellappa. A dual-path model with adaptive attention for vehicle re-identification. In *ICCV*’19
- [22] D. P. Kingma and J. Ba. Adam: A Method for Stochastic Optimization. In *ICLR*, 2015.
- [23] R. Kumar, G. Charpiat, and M. Thonnat. Multiple Object Tracking by Efficient Graph Partitioning. In *ACCV*’14
- [24] R. Kumar, E. Weill, F. Aghdasi, and P. Sriram. Vehicle re-identification: an efficient baseline using triplet embedding. In *IJCNN*’19
- [25] L. Liao, R. Hu, J. Xiao, Q. Wang, J. Xiao, and J. Chen. Exploiting effects of parts in fine-grained categorization of vehicles. In *ICIP*, 2015.
- [26] S. Liao, Y. Hu, X. Zhu, and S. Z. Li. Person re-identification by Local Maximal Occurrence representation and metric learning. In *CVPR*, 2015.
- [27] Y. L. Lin, V. I. Morariu, W. Hsu, and L. S. Davis. Jointly optimizing 3D model fitting and fine-grained classification. In *ECCV*, 2014.
- [28] H. Liu, Y. Tian, Y. Wang, L. Pang, and T. Huang. Deep Relative Distance Learning: Tell the Difference Between Similar Vehicles. In *CVPR*, 2016.
- [29] X. Liu, W. Liu, H. Ma, and H. Fu. Large-scale vehicle re-identification in urban surveillance videos. *ICME*, 2016.
- [30] X. Liu, H. Ma, H. Fu, and M. Zhou. Vehicle Retrieval and Trajectory Inference in Urban Traffic Surveillance Scene. In *ICDSC*, 2014.
- [31] Y. Lou, Y. Bai, J. Liu, S. Wang, and L.-Y. Duan. A large-scale dataset for vehicle re-identification in the wild. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2019.
- [32] P. Luo, C. C. Loy, X. Tang, L. Yang, P. Luo, C. C. Loy, and X. Tang. A Large-Scale Car Dataset for Fine-Grained Categorization and Verification. In *CVPR*, 2015.

- [33] B. Ma, Y. Su, F. Jurie, B. Ma, Y. Su, and F. Jurie. Local Descriptors Encoded by Fisher Vectors for Person Re-identification. In *ECCV Workshops*, 2012.
- [34] B. P. Ma, Y. Su, and F. Jurie. BiCov: a novel image representation for person re-identification and face verification. In *BMVC*, 2012.
- [35] R. Manmatha, C. Y. Wu, A. J. Smola, and P. Krahenbuhl. Sampling Matters in Deep Embedding Learning. In *CVPR*, 2017.
- [36] A. Mishchuk, D. Mishkin, F. Radenovic, and J. Matas. Working hard to know your neighbor’s margins: Local descriptor learning loss. In *NIPS*, 2017.
- [37] V. Mnih, N. Heess, A. Graves, and K. Kavukcuoglu. Recurrent Models of Visual Attention. In *NIPS*, 2014.
- [38] M. Naphade, M.-C. Chang, A. Sharma, C. Anastasiu, David, V. Jagarlamudi, P. Chakraborty, T. Huang, S. Wang, M. Y. Liu, R. Chellappa, J.-N. Hwang, and S. Lyu. The 2018 NVIDIA AI City Challenge. *CVPR Workshops*, 2018.
- [39] O. Rippel, M. Paluri, P. Dollar, and L. Bourdev. Metric Learning with Adaptive Density Discrimination. In *ICLR*, 2016.
- [40] E. Ristani and C. Tomasi. Features for Multi-Target Multi-Camera Tracking and Re-Identification. In *CVPR*, 2018.
- [41] F. Schroff, D. Kalenichenko, and J. Philbin. FaceNet: A unified embedding for face recognition and clustering. In *CVPR*, 2015.
- [42] L. Shen, Z. Lin, and Q. Huang. Relay Backpropagation for Effective Learning of Deep Convolutional Neural Networks. In *ECCV*, 2016.
- [43] Y. Shen, T. Xiao, H. Li, S. Yi, and X. Wang. Learning Deep Neural Networks for Vehicle Re-ID with Visual-spatio-Temporal Path Proposals. *ICCV*, 2017.
- [44] K. Simonyan and A. Zisserman. Very deep convolutional networks for large-scale image recognition. *CoRR*’14
- [45] K. Simonyan and A. Zisserman. Very Deep Convolutional Networks for Large-Scale Image Recognition. In *ICLR*, 2015.
- [46] J. Sochor, A. Herout, and J. Havel. BoxCars: 3D Boxes as CNN Input for Improved Fine-Grained Vehicle Recognition. In *CVPR*, 2016.
- [47] J. Spanhel, J. Sochor, R. Juranek, A. Herout, L. Marsik, and P. Zemcik. Holistic recognition of low quality license plates by CNN using track annotated data. *AVSS*, 2017.
- [48] S. Tang, M. Andriluka, B. Andres, and B. Schiele. Multiple people tracking by lifted multicut and person re-identification. In *CVPR*, 2017.
- [49] Z. Tang, M. Naphade, S. Birchfield, J. Tremblay, W. Hodge, R. Kumar, S. Wang, and X. Yang. Pamtri: Pose-aware multi-task learning for vehicle re-identification using highly randomized synthetic data. In *ICCV*’19
- [50] Z. Tang, M. Naphade, M.-Y. Liu, X. Yang, S. Birchfield, S. Wang, R. Kumar, D. Anastasiu, and J.-N. Hwang. Cityflow: A city-scale benchmark for multi-target multi-camera vehicle tracking and re-identification. In *CVPR*’19
- [51] O. Tuzel, F. Porikli, and P. Meer. Region covariance: A fast descriptor for detection and classification. In *European Conference on Computer Vision*, pages 589–600, 2006.
- [52] Y. Wang, L. Xie, S. Qiao, Y. Zhang, W. Zhang, and A. L. Yuille. A Deep Learning-Based Approach to Progressive Vehicle Re-identification for Urban Surveillance. In *ECCV*, 2016.
- [53] Z. Wang, L. Tang, X. Liu, Z. Yao, S. Yi, J. Shao, J. Yan, S. Wang, H. Li, and X. Wang. Orientation Invariant Feature Embedding and Spatial Temporal Regularization for Vehicle Re-identification. In *ICCV*, 2017.
- [54] K. Q. Weinberger and L. K. Saul. Distance Metric Learning for Large Margin Nearest Neighbor Classification. *The Journal of Machine Learning Research*, 10:207–244, 2009.
- [55] L. Wen, D. Du, Z. Cai, Z. Lei, M. Chang, H. Qi, J. Lim, M. Yang, and S. Lyu. UA-DETRAC: A new benchmark and protocol for multi-object detection and tracking.
- [56] N. Wojke and A. Bewley. Deep Cosine Metric Learning for Person Re-identification. In *WACV*, 2018.
- [57] T. Xiao, H. Li, W. Ouyang, and X. Wang. Learning Deep Feature Representations with Domain Guided Dropout for Person Re-identification. In *CVPR*, 2016.
- [58] K. Yan, Y. Tian, Y. Wang, W. Zeng, and T. Huang. Exploiting Multi-grain Ranking Constraints for Precisely Searching Visually-similar Vehicles. In *ICCV*, 2017.
- [59] D. Zapletal, A. Herout, and A. Herout. Vehicle Re-Identification for Automatic Video Traffic Surveillance. In *CVPR Workshops*, 2016.
- [60] L. Zhang, T. Xiang, and S. Gong. Learning a Discriminative Null Space for Person Re-identification. In *CVPR*, 2016.

[61] Y. Zhou and L. Shao. Vehicle Re-Identification by Adversarial Bi-Directional LSTM Network. In WACV, 2018.



**Ratnesh Kumar** is currently a Deep Learning Architect at NVIDIA. He has obtained his Ph.D. from STARS team at Inria, France in 2014. During his Ph.D., he was focused on long term video segmentation using optical flow and multiple object tracking. Subsequently he worked as Postdoc at Mitsubishi Electric Research Labs (MERL) Cambridge, Boston, on detecting actions in streaming videos. He also holds Bachelors in Engineering from Manipal University, India and Master of Science from University of Florida at Gainesville, USA. At NVIDIA, he works on deep learning solutions to problems in video analytics ranging from object detection to reidentification and action detection.



**Eddie Weill** is a Senior Data Scientist at NVIDIA specializing in assisting enterprise customers in the development of deep learning, machine learning, and accelerated computing solutions. Prior to his current position, he worked at NVIDIA in the Intelligent Video Analytics group developing deep learning models for smart cities. As part 22 of his Ph.D., he focused on deep learning for com-

[62] Y. Zhou and L. Shao. Viewpoint-aware Attentive Multi-view Inference for Vehicle Re-identification. In CVPR, 2018.

puter vision and its applications to embedded systems. During his graduate studies at Clemson University, he focused on deep learning as well as highperformance computing applications.



**Farzin Aghdasi** is a senior software manager with NVIDIA. He holds a Ph.D. from the University of British Columbia in Electrical Engineering. He was previously a group master expert at Schneider Electric, and a tenured associate professor at the university of the Witwatersrand. His research interests are in areas of image and video processing and deep learning architectures for intelligent video analytics. He has published over 50 articles in refereed international journals and conference proceedings, and has over 30 granted patents.



**Partha Sriram** is a Chief Systems Software Architect for Video Analytics Platforms & V.P. Of Engineering at NVIDIA. He holds a Ph.D. from University of Arizona, USA.