

Rethinking Person Re-Identification with Confidence

George Adaimi
EPFL

george.adaimi@epfl.ch

Sven Kreiss
EPFL

sven.kreiss@epfl.ch

Alexandre Alahi
EPFL

alexandre.alahi@epfl.ch

Abstract

A common challenge in person re-identification systems is to differentiate people with very similar appearances. The current learning frameworks based on cross-entropy minimization are not suited for this challenge. To tackle this issue, we propose to modify the cross-entropy loss and model confidence in the representation learning framework using three methods: label smoothing, confidence penalty, and deep variational information bottleneck. A key property of our approach is the fact that we do not make use of any hand-crafted human characteristics but rather focus our attention on the learning supervision. Although methods modeling confidence did not show significant improvements on other computer vision tasks such as object classification, we are able to show their notable effect on the task of re-identifying people outperforming state-of-the-art methods on 3 publicly available datasets. Our analysis and experiments not only offer insights into the problems that person re-id suffers from, but also provide a simple and straightforward recipe to tackle this issue.

1. Introduction

According to the nineteenth-century physicist James Clerk Maxwell, doubt, which he describes as "thoroughly conscious ignorance", is the prelude to science. Doubt allows us to question our decision and pushes us to thoroughly find a detailed reason behind them. In many perception tasks, state-of-the-art neural networks rarely include the notion of doubt while training. They are trained to output a high probability for the correct class. Thus, when dealing with images of different classes but very similar characteristics, they would focus on relatively unimportant variations to distinguish the different inputs. This would lead to unreasonable classification of the challenging images in an effort to reduce its loss function. This describes an inherent problem in the popular person re-identification task. In the same way humans can be doubtful of their own decisions, a model should be allowed to have some doubt in its classification.

Person re-identification methods attempt to extract dis-



Figure 1: In person re-identification systems, visually similar individuals are difficult to discern thus forcing the model to focus on unnecessary variations. We propose reducing a model’s confidence as a solution for this problem. Our model with penalized confidence correctly ranks the gallery images (top-3, left to right). However, the baseline model (top) focuses on the pose (first image) and the shirt (second image) leading to incorrect identification. Red frame indicates wrong ID while green frame indicates correct ID compared to the query. Best viewed in color.

criminative features from two images and measure how similar these extracted features are. In addition to the use of different metric learning methods, a key milestone in person re-identification is the use of cross-entropy to learn representations that are distinct for different identities[35, 36]. Since then, an arms race of methods built on top of this by making use of different human-specific characteristics (e.g., human semantic segmentation, pose). A main pitfall of learning a representation with cross-entropy supervision is the fact that it separates the different inputs solely based on the labels without taking into consideration the actual similarity between the inputs. None of the recent methods have tackled the problem where, even though two very similar individuals are distinct, their similarity score should encode information about how similar they appear while also distinguishing them. The network usually tries its best to

find a boundary between the different classes even for inputs that are very close together. This leads the network to find unreasonable explanations for the differences in labels and thus would negatively affect its parameters. After long enough training, the network would increasingly become more confident about its decision. Since person re-id also deals with the problem of having a small set of images per class, it would aggravate this issue. Controlling the network’s confidence about its prediction would alleviate this problem (Figure 1). Even though this concept has been studied before, it has not been studied in the field of person re-identification.

In this paper, we propose to model confidence when learning representations appropriate for person re-id. By inducing doubt while training a network, we are able to tackle the inherent problem discussed previously when cross-entropy is used in a distance metric and representation learning problem such as person re-identification. Inspired by previous works that use uncertainty to regularize the network, we study three alternatives that aim at reducing the confidence of the network and show a gain of 6-7 % in mAP across 3 different datasets. Although these methods haven shown only a small improvement on other image classification tasks [28, 2, 18], they drastically improve the performance of person re-id models due to its innate problem (Section 3). By combining our methods with advanced ranking methods, we outperform state-of-the-art models without modeling characteristics specific to humans. The software is open source and available online.¹

2. Related Work

With the prevalence of deep neural networks in most computer vision tasks, person re-id followed this success when Li et al. [15] introduced a deep learning method for re-id that tried to overcome the problems of bounding box misalignment, photometric and geometric transforms while also introducing a new bigger dataset specifically for this task. This paved the way for new methods and datasets to emerge, causing the person re-id performance of machines to improve. Other work developed new methods that tackle specific challenges in person re-id by introducing different architectures and modules [14, 16, 26, 33].

Attention in Person reID. Recently, many methods have tried to improve the representation of the input by training multiple networks that extract global and local features and then combine these features to form the final representation. This is usually done by using either a deterministic way of dividing the different parts of the representation [41, 6, 31, 1] or making use of attention modules to separate the different parts [16, 43, 38]. Other works extracted intermediate representations to gain information about the

input at different levels arguing that this allows the network to learn distinctive characteristics of the input at different scales [3, 32, 33]. Even though these methods showed improvement over their predecessors, these methods usually require separate networks to process each of the different features, leading to a more complex architecture and training procedure.

Human Characteristics. Another direction other researchers have taken is to make use of information and characteristics related specifically to humans in order to improve person re-id. The work by Xu et al. [37] aims at detecting three different types of pose information such as keypoints, rigid body parts (e.g., torso), and non-rigid parts (e.g., limbs). These information were extracted using an off-the-shelf human pose estimator. Then, with the help of these body parts, the features extracted by a feature extractor are refined and used to classify the different individuals. The use of third-party methods make their model highly dependent on the performance of these methods. Another approach by Sarfraz et al [21] uses keypoint information, in addition to the input image, to train a ResNet-50 model as well as another connected module that detects the view (front, back or side). Kayaleh et al. [11] also made use of features extracted from different body parts and concatenated them to form a global feature which in turn was used to perform re-identification. The disadvantage of these methods is their high dependence on other methods and datasets that require annotation. Moreover, the fact that these models depend on specific human characteristics prevents them from being leveraged for other image-retrieval and clustering tasks.

Re-Ranking. In addition to learning better features, many works have tried to improve the ranking process of person re-id by including information about how the different galleries are related instead of just using the relationship between the pairs of queries and galleries [46, 39, 24, 46, 8, 13, 39, 40]. Zhong et al [46] introduced a method for refining the distances between the queries and galleries by making use of the k-reciprocal nearest neighbors. This is done as a post-processing step to improve the ranking process. Shen et al. [24] argued that this does not help in learning better features during training and introduced a new learnable module that performs a random walk on a graph connecting the different gallery images. By performing a random walk operation, gallery-to-gallery (G2G) information is taken into consideration while training the network, thus resulting in a more complete representation that provides a better ranking performance. Other methods also tried to include G2G information by using Graph Neural Networks [23] and Conditional Random Fields [4]. We will make use of G2G information by applying different re-ranking methods.

Metric Learning. Several previous works have tried to

¹Link will be revealed after review process



Figure 2: Pairs of images of different IDs but very similar appearance - Market1501.

tackle the problem of person re-id by introducing new metric loss functions. Both contrastive [9] and binary loss functions have been employed in order to push apart negative image pairs while pulling positive image pairs together [19, 1]. Taking into consideration both the pull and push of contrastive loss, other methods [33, 17, 32] used triplet loss that simultaneously tackles negative and positive pairs leading to a less greedy method. Chen et al. [5] extended this loss to quadruplet inputs. The drawback of these methods is their high sensitivity to the sampling technique used. As a result, Yu et al. [42] introduced the HAP2S loss to tackle this drawback and showed improvement in performance. All the above methods try to encode metric information in the embedding space compared to cross-entropy which is considered as a representation learning method.

In this paper, we do not make use of human characteristics or feature division and show the importance of confidence when training a person re-id model with a cross-entropy loss.

3. Problem Formulation

A person re-id model’s main task is to distinguish between people across frames. As previously stated, the person re-identification task is a challenging task since it tries to relate images of people across different cameras. The fact that the images are captured by different cameras might lead to subtle differences in hue and image color that can drastically effect the performance of a re-id model. Moreover, the illumination, background clutter, occlusion, observable human body parts, and perceived posture of the person are usually dramatically different which might easily fool the network and render it unusable. Even images of people captured by the same camera can have many of these variations.

Due to the challenges explained above, there isn’t always a clear margin of separation between individuals. People in some cases have very subtle differences that separate them from each other making the task of identifying them even more challenging for a human observer. A good example is shown in Figure 2 which introduces the inherent challenge we are trying to tackle in this paper.

The people within the images in Figure 2 are very difficult discern from one another even for a human eye. Each

pair of images shows two different people who share very similar appearances. When a model is trained to separate these images, it might face difficulties doing so. In order to reduce its loss in this case, the network will learn to focus on pose or even the illumination of the image. These two variations are some of the many variations that previous methods are trying to overcome. Current state-of-the-art person re-id systems train their own models by using the cross-entropy loss function. The cross-entropy calculates the number of bits needed for an event, which in this case is the label given the input, using the estimated probability distribution instead of the true distribution. In the case of training a neural network, the cross-entropy is minimized so that the model distribution is the same as that of the ground-truth, which is a one-hot encoding in person-reid. This means minimizing this loss pushes the distribution of the model to output a high probability for the correct label while outputting very low probabilities for the others. The fact that cross-entropy requires that the logits for the ground-truth label to be much bigger than other labels pushes the network to take into consideration certain destructive variations to separate the different classes and especially for images such as in Figure 2.

In order to modify the cross entropy in a way that solves the problem described above, we add a missing term to the loss function which allows it to not be confident about certain datapoints. Thus, the modified cross entropy loss function allows the network not to overfit on variations that are destructive for the person re-id task and accept the fact that people do sometimes look very similar. The idea of preventing the network from being very confident is not a new concept. However, its evaluation on other computer vision tasks only lead to slight improvements in performance. From the reasoning based on Figure 2 as well as the characteristics of person re-id datasets, we show in this paper that this concept, if applied to a simple baseline, can improve the results drastically and even outperform certain highly specialized state-of-the-art methods.

4. Method

To reiterate, current person re-id models face difficulties in distinguishing between different individuals who share some visual similarities due to the model’s objective of maximizing its confidence in its predictions. In this section, we introduce three different methods that allow the network to be less confident about the different labels. These methods usually show a small improvement when used during training in other computer vision tasks [29, 18, 2]; however, we show that, because of the problems specified in Section 3, these methods provide a drastic boost for the task of person re-id.

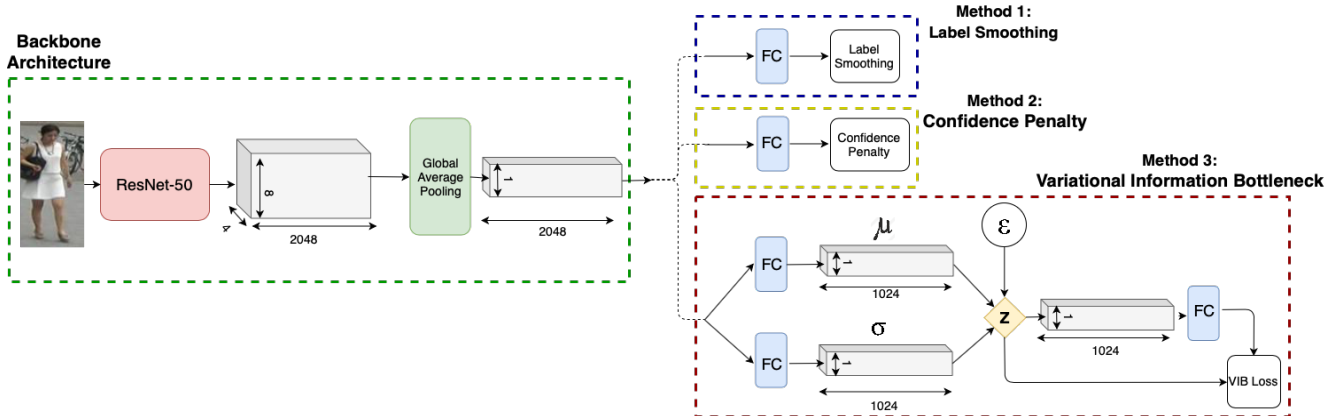


Figure 3: Network architecture including the three methods being studied: Label Smoothing (LS), Confidence Penalty (CP), Variational Information Bottleneck (VIB). ϵ is the Gaussian noise needed for the reparametrization trick, μ and σ are the mean and standard deviation respectively of the latent Z distribution.

4.1. Label Smoothing

Label smoothing is a form of model regularizer introduced by Szegedy et al. [29] which aims at allowing a model to be less confident about a certain prediction. It regularizes a softmax classifier by assigning a small value to all ground-truth labels. This is done by changing the ground-truth distribution ($q(c|x)$) that the model is trying to approximate to a smoother distribution ($q_{LS}(c|x)$):

$$q_{LS}(c|x) = \begin{cases} 1 - \frac{(C-1)\epsilon}{C} & c = \text{label}(x), \\ \frac{\epsilon}{C} & \text{otherwise.} \end{cases} \quad (1)$$

This method makes sure that the label for the correct class does not become much larger than all other classes and thus prevents the network from overfitting. When label smoothing was proposed and tested on ImageNet, it showed a small improvement of around 0.2% for top-1 error. Even though it did not show a huge improvement, we show in Section 6 that this method has a bigger effect on the task at hand based on the arguments stated in Section 3. As can be observed in Figure 3, this method requires only a modification to the cross-entropy loss function where the modified ground-truth distribution is used and the resulting loss function would be:

$$L_{LS} = \alpha H(q_{LS}(c|x); p(\mathbf{c}|\mathbf{x})), \quad (2)$$

where $p(c|x)$ is the model's output distribution.

4.2. Confidence Penalty

While the network trains, its predictions become more and more confident, giving more probability to a specific class compared to other classes. Having confident predictions indicates that the output distribution $p(\mathbf{c}|\mathbf{x})$ over all

the classes \mathbf{c} has low entropy since one label dominates the prediction. Its entropy can be calculated by:

$$H(p(\mathbf{c}|\mathbf{x})) = - \sum_i p(\mathbf{c}_i|\mathbf{x}) \log(p(\mathbf{c}_i|\mathbf{x})).$$

This equation measures the uncertainty of a model in performing its prediction. In order to make the network less certain, Pereyra et al. [18] suggested penalizing the entropy of the output distribution. They showed that by doing so, they got a smoother output distribution as well as a small improvement on MNIST. This method however did not show an improvement on a more difficult dataset such as CIFAR-10. By penalizing the entropy, the loss function becomes:

$$L_{CP} = \alpha L_{cross} - \beta H(p(c|x)), \quad (3)$$

where β controls how much to penalize the $H(p(\mathbf{c}|\mathbf{x}))$ and α controls the strength of the cross-entropy loss (L_{cross}). This method is similar to label smoothing in that it allows the network to output a small probability to labels different than the ground-truth. Similar to label smoothing, no architecture modification is required except adding the confidence penalty to the loss (Equation 3). This is shown in Figure 3.

4.3. Deep Variational Information Bottleneck

The information bottleneck (IB) principle [30] is a technique that tries to find the best trade-off between accuracy and complexity of latent variables. Latent variables are hidden variables that describe a specific input while maintaining all the relevant information needed for a specific task. The information bottleneck method tries to minimize this objective:

$$\min_{p(z|x)} I(X; Z) - \beta I(Z; Y), \quad (4)$$

where Z is the latent variable, X is the input, and Y is the output. Based on the above equation, the objective is to learn a representation Z that is very informative about Y while compressive about X . In order to apply the IB objective to a neural network, Alemi et al. [2] approximated a lower bound to the information bottleneck by using variational inference and the reparameterization trick introduced by Kingma et al. [12] to introduce a new objective function referred to as Variational Information Bottleneck (VIB).

When applying this method, the model is divided into an encoder that takes the input X and maps it to a distribution describing the latent space Z . The encoder outputs both the mean μ and standard deviation σ that describe this distribution. Then the predicted distribution is used to sample a specific latent representation. To force the first part of equation 4 to be minimized, this distribution should not depend on the input thus forcing the representation Z to forget some information about it. This is done by minimizing the divergence between the encoder’s distribution $p(z|x)$ and the prior $r(z)$. The resulting objective function that is minimized:

$$L_{VIB} = \alpha L_{cross} + \beta KL[p(\mathbf{Z}|x), r(\mathbf{Z})]. \quad (5)$$

In order to compute the KL divergence analytically and backpropagate using its gradients, $p(\mathbf{Z}|x)$ is approximated by a multivariate Gaussian distribution with a diagonal covariance matrix while $r(\mathbf{Z})$ is an isotropic multivariate Gaussian. As can be seen in equation 5, if $\beta \rightarrow \infty$, the latent representation would follow a distribution independent of the input and thus different classes will have similar representations. This is somewhat similar to the effect of both confidence penalty and label smoothing where a single representation is forced to contain some information about more than one label. However, VIB applies this restriction directly to the latent space. Using this method while training, Alemi et al. [2] showed close results to state-of-the-art models while using less information about the input which is measured using mutual information $I(X; Z)$. Compared to previously mentioned methods, in order to use the VIB loss, a fully connected layer is added at the output of the ResNet-50 base model to compute the mean and standard deviation as shown in Figure 3.

4.4. Methods Comparison

Analyzing all three methods reveals the similar effect that they share together as well as their differences. These methods aim at increasing uncertainty in the training procedure of the model. Label smoothing(LS) and confidence penalty (CP) achieve this by making sure that the network is less penalized on wrong classifications. VIB pushes the representations to be more independent of the input and label. Moreover, both confidence penalty and label smoothing can be expressed by a KL divergence between the output and a

uniform distribution. The difference, however, is that label smoothing can be expressed as $KL[u||p(\mathbf{y}|x)]$ while confidence penalty can be expressed as $KL[p(\mathbf{y}|x)||u]$ where u represents the uniform distribution. This difference has a significant effect on the training and the representations learned since in label smoothing, the error is weighted by the uniform distribution ($\frac{1}{N_c}$ where N_c is the number of classes). On the contrary, using confidence penalty weighs the error by the output distribution itself. In other words, when confidence penalty is used, the divergence between the output distribution and the uniform u is affected by the network’s current confidence about the input compared to label smoothing, which is mainly affected by u . In summary, all three methods can be expressed as a KL divergence where label smoothing and confidence penalty act on the output while VIB acts directly on the representation.

5. Experiments

To evaluate our proposed method, we use three publicly available person re-identification datasets which are Market-1501 [44], the recently created dataset MSMT17 [34], and DukeMTMC-reID [45].

Market1501: The Market dataset is a well-known person re-identification dataset that contains 32,668 bounding boxes of 1,501 individuals captured using 6 cameras. These bounding boxes were obtained using the Deformable Part Model (DPM) [7]. The training set is made up of 751 identities with 12,936 images while the test set has 750 identities distinct from the one in the training set divided into query and gallery images.

MSMT17: This is a very recent dataset which was carried out over a long period of time. This benchmark contains a total of 126,441 bounding boxes of 4,101 identities captured using 15 cameras. The images vary in terms of location (outdoors, indoors), weather condition (over a month), as well as different times of day (morning, noon, afternoon). The bounding boxes were obtained using Faster RCNN and corrected using labelers. Containing so much variations makes this dataset challenging as well as a good benchmark to use.

DukeMTMC-reID: The DukeMTMC-reID dataset is a small part of the bigger DukeMTMC dataset that is usually used for multi-target multi-camera tracking. It is taken from 8 different cameras, and the person bounding box is manually labeled. It is made up of 1,404 different identities with 702 identities used for training and 702 other identities used for testing.

5.1. Evaluation Protocol

For evaluation, we use the cumulative matching characteristic (CMC) and mean Average Precision (mAP). These two metrics are the most popular evaluation metrics since

Parameters	Market-1501	DukeMTMC	MSMT17
LR_{cross}, α	$5 \times 10^{-4}, 2$	$2 \times 10^{-4}, 1$	$3 \times 10^{-4}, 4$
LR_{LS}, α	$5 \times 10^{-4}, 2$	$5 \times 10^{-4}, 5$	$3 \times 10^{-4}, 5$
LR_{CP}, α	$6 \times 10^{-4}, 3$	$6 \times 10^{-4}, 3$	$4 \times 10^{-4}, 5$
LR_{VIB}, α	$4 \times 10^{-4}, 6$	$6 \times 10^{-4}, 3$	$5 \times 10^{-4}, 6$
β_{CP}	0.085	0.085	0.085
β_{VIB}	0.01	0.01	0.01

Table 1: Hyperparameters for the different datasets and methods. LR: learning rate, β : pre-factors for loss constraint, α : pre-factor for cross-entropy

person re-identification systems should be able to output all the correct matches (mAP) in addition to having high accuracy at different ranks (CMC). During testing, for every query there is a list of gallery images ordered in increasing order according to their L2 distance from this query.

5.2. Implementation Details

The model was pre-trained using ImageNet. We do not add any layer to ResNet-50 when training both using label smoothing and confidence penalty except for a fully connected layer that outputs the different labels. When training the VIB algorithm, a fully-connected layer was added before the classification layer to output the mean and standard deviation which describe the distribution of the latent representations. A latent variable is then sampled from the predicted distribution. For all methods and datasets, hyperparameter tuning was performed for ResNet-50 in order to get the best possible accuracy.

Data Augmentation. We follow methods of data augmentations that are commonly used in the field of person re-identification. Since Market1501 uses DPM to obtain the bounding boxes, the images are initially randomly cropped. For all datasets, the inputs are resized to 256x128. Before providing them to the network, a random rectangle, with pixel values randomly assigned between [0, 255], is erased [47] from the images, and the resulting images are flipped horizontally with a probability of 0.5. This makes the network more robust to the orientation of the people in the image as well as occlusion. Each image is then normalized and standardized using the mean and standard deviation provided when using a model pretrained on ImageNet. These transformations were applied only for the training set.

Hyperparameter Tuning. Since the hyperparameters (e.g., learning rate, β , and α) we are trying to optimize have multiplicative effects on the training procedure, the best method is to perform a log-space search. This is due to two reasons. The parameter is not too sensitive such that there may not be too much difference with 10 and 15 compared to 10 and 1000. The other reason is that using logarithmic

scales allows us to search over a bigger space quickly.

Training Procedure. The samples used to form the training batch are randomly sampled from the datasets. It does not require any special sampling such as the PK Sampling required by triplet loss[22], which randomly samples P identities and then randomly K images for each identity to form a batch. The mini-batch has a size of 32 images. The model is trained for 300 epochs using the AMSGrad [20] optimizer for all datasets with the learning rate decaying by 10 at epoch 20 and 40. In order to make sure that all models were trained with the best parameters, we perform hyperparameter tuning, as discussed previously. The different hyperparameters for the different datasets are shown in Table 1.

Evaluation Procedure. For testing, the features that are extracted just before the last classification layer are used for the ranking process. The features for the queries and galleries are extracted and then compared to rank the gallery images relative to each query image. This is done when label smoothing or confidence penalty is used. When using the VIB loss, the network has an additional fully connected layer that outputs the mean and standard deviation for every latent dimension and a reparameterization trick that depends on a random Gaussian noise. For ranking, we use the mean produced by the model as features for each image since this represents the average of the distribution over which the input image is mapped to. This is also due to the fact that the standard deviations tend to 1. To the best of our knowledge, using a latent representation, sampled from a Gaussian parametrized by the predict mean and standard deviation, has not been tackled before for the person re-id task.

6. Results

In order to show both qualitative and quantitative results, we split our results into three parts. In Sections 6.1 and 6.2, we compare our proposed methods to published baseline results and state-of-the-art methods respectively. In Section 6.3, we investigate the effect of the three methods on the ranking process of person re-id. Although these methods were tested on ResNet-50, other re-id models can benefit from their positive effect on the performance especially when dealing with visually similar individuals.

6.1. Properly Trained Baseline

We compare our baseline to previously reported results of ResNet-50 on the Market-1501 and DukeMTMC-reID datasets. The published results reported in Table 2 correspond to pre-trained ResNet-50 that used the cross-entropy loss similar to our method. As can be observed in Table 2, there is a clear difference between our result and the results reported in published papers as well as amongst the published results themselves. Our properly trained base-

Model	Market1501		DukeMTMC	
	mAP	Rank1	mAP	Rank1
ResNet-50 [17]	47.78	73.90	44.99	65.22
ResNet-50 [24]	59.8	81.4	55.5	75.3
ResNet-50 [21]	59.8	82.6	50.3	71.5
ResNet-50 [3]	66.0	84.3	48.6	71.6
ResNet-50 [10]	66.95	84.42	57.34	75.60
ResNet-50 [11]	66.32	85.10	54.77	73.70
Our ResNet-50	70.2	87.5	59.6	78.6

Table 2: Comparison with published ResNet-50 results on the Market-1501 and DukeMTM-reID dataset.

line, which consists of a ResNet-50 model trained using a normal cross-entropy loss, was able to outperform all previous baselines. This table represents one of the many pitfalls that occurs when training a model. This is shown by the fact that papers that make use of exactly the same baseline have different results. This is usually due to the hyperparameters chosen. Another pitfall is that to compare different baselines and losses, the same hyperparameters are set. This is somewhat unfair since different baselines and losses optimize different parameters and in different ways thus requiring distinct hyperparameters. This is why we employ different learning rates for different datasets and methods as shown in Table 1. As a result, we were able to achieve, using the baseline, around $\sim 3\%$ increase in mAP and rank-1 for both datasets.

6.2. Comparison with State-of-the-art

We evaluate our proposed confidence-based methods against recently published papers in person re-id. Each of our methods are evaluated on three datasets: Market1501, DukeMTMC-reID, and MSMT17. We are able to reach state-of-the-art performance without any human specific design and added complexity thus showing the importance of penalizing the confidence of a network in person re-id. We also do not make use of data augmentation during the evaluation stage like DuATM [25].

Evaluation on Market1501: As shown in Table 3, the models were able to reach state-of-the-art results. In order to better understand the importance of penalizing confidence compared to other methods, it is important to note some distinct differences. Confidence penalty was able to outperform HAP2S [42] which tried to deal with hard samples by giving them higher weights. Moreover, Mancs[32], which shows good performance, makes use of three different losses, attention layers, as well as a special sampling scheme. To compare our results with methods that include gallery-to-gallery information during inference, such as Deep Group RW [24] and SGGNN [23], we apply re-ranking to our three methods. We were able to outperform

Model	Market1501		DukeMTMC	
	mAP	Rank1	mAP	Rank1
CamStyle (R)[48]	71.55	89.49	57.61	78.32
HAP2S_E+Xent(R)[42]	74.49	89.73	62.62	79.08
DuATM(!R)[25]	75.22	89.96	63.14	81.46
MLFN (!R)[3]	74.3	90.0	62.8	81.0
Shen et al.(R)[24]	75.3	90.1	63.2	80.3
PSE(R)[21]+ECN	84.0	90.3	79.8	85.2
DaRe(!R)[33]+RR	86.7	90.9	80.0	84.4
SPReID ^{w/fg} (!R)[11]*	78.66	90.97	65.66	81.73
HA-CNN (!R) [16]	75.7	91.2	63.8	80.5
DuATM(!R)[25]**	76.62	91.42	64.58	81.82
SPReID ^{comb} (!R)[11]*	79.67	91.45	68.78	83.3
P-Aligned (!R)[27]	79.6	91.7	69.3	84.4
SGGNN(R)[23]	82.8	92.3	68.2	81.1
Deep Group RW(R)[24]	82.5	92.7	66.4	80.7
Mancs(R)[32]	82.3	93.1	71.8	84.9
DNN+CRF(R) [4]	81.6	93.5	69.5	84.9
P-Aligned (!R)[27]+RR	89.9	93.4	83.9	88.3
Our ResNet	70.7	87.2	59.6	78.6
Our ResNet(VIB)	76.1	90.2	62.4	80.7
Our ResNet(LS)	76.7	91.0	64.4	82.7
Our ResNet(CP)	78.2	91.4	66.8	83.9
Our ResNet+RR	85.7	89.7	78.5	83.4
Our ResNet(VIB)+RR	88.6	91.8	79.0	84.3
Our ResNet(LS)+RR	89.1	92.2	82.2	86.6
Our ResNet(CP)+RR	90.0	92.6	83.5	87.4
Our ResNet(VIB)+ECN	88.2	92.0	78.9	85.1
Our ResNet(LS)+ECN	89.4	92.7	83.2	86.9
Our ResNet(CP)+ECN	90.1	93.1	84.1	88.5

Table 3: Comparison with state-of-the-art methods on Market-1501 and DukeMTMC-reID. (!R): uses model different than ResNet, (R): uses ResNet-50, ECN: Expanded Cross Neighborhood Re-Ranking[21], "RR": k-reciprocal re-ranking[46], Xent: Softmax, *: uses combination of 10 datasets for training, **: uses data augmentation during evaluation stage.

these methods with a significant increase in mAP($\sim 7.5\%$). As a result, we got state-of-the-art performance without the added complexity of learning new layers and parameters while tackling the problem stated in Section 3.

Evaluation on DukeMTMC-reID: Similar to the Market-1501 dataset, we achieved competitive results in all proposed methods with confidence penalty resulting in the best improvement (Table 3). In addition to that, using Sarfraz et al.'s [21] recent re-ranking method (ECN), we were able to get better results than PSE [21] in both mAP and rank-1. It is important to note that SPReID augments the training data of both DukeMTMC-reID and Market1501 with 10 datasets resulting in a large amount of training samples which would

MSMT17			
Model	mAP	Rank1	Rank10
GoogleNet[34]	23.0	47.6	71.8
PDC[34]	29.7	58.0	79.4
GLAD[34]	34.0	61.4	81.6
Our ResNet	31.8	59.3	80.2
ResNet-50(VIB)	35.1	66.2	84.1
ResNet-50(LS)	36.9	66.8	84.9
ResNet-50(CP)	39.3	68.6	85.3
Our ResNet + RR	49.8	65.7	79.8
Our ResNet(VIB) +RR	55.4	73.3	84.7
Our ResNet(LS) + RR	57.1	73.7	85.3
Our ResNet(CP) + RR	59.1	75.3	85.8

Table 4: Comparison with state-of-the art on the MSMT17 dataset.

improve the performance of the network.

Evaluation on MSMT17: Since this is a bigger dataset with many variations, it proved to be a challenging benchmark[34]. Nonetheless, we were able to show a notable improvement over previous methods as well as over our own baseline (Table 4). Similarly, confidence penalty performed the best by achieving 68.6% in rank-1 and 39.3% in mAP. By applying re-ranking, both rank-1 and mAP are further improved to 75.3% and 59.1% respectively.

6.3. Effect of Proposed Methods

In addition to achieving state-of-the art performance, it is also important to understand the effect of these three methods on the ranking process. All three methods aim at allowing the network to share some representation among different classes. This prevents the network from focusing on undesirable information when separating very similar looking individuals. To show this effect, we compare the confidence penalty model against the baseline model since it resulted in the best performance (Figure 4). As can be seen, the test samples presented in Figure 4 are difficult to rank even for a human observer. This confirms the intrinsic difficulty of person re-id stated in Section 3. When confidence penalty is not used for training, the network focuses on unimportant variations between the images. For instance, in both sets of samples, the incorrect gallery images are very similar to the query image despite belonging to a different person. The baseline links the query image to the gallery images by possibly focusing on the background, shirt color, posture, and body rotation of the individual in question. These characteristics are typically features that can confuse the model leading to wrong identification. Adding the confidence penalty is observed to remedy this challenge, as can be seen for all test samples provided. Adding the confidence penalty helps the model capture the subtle differences between multiple

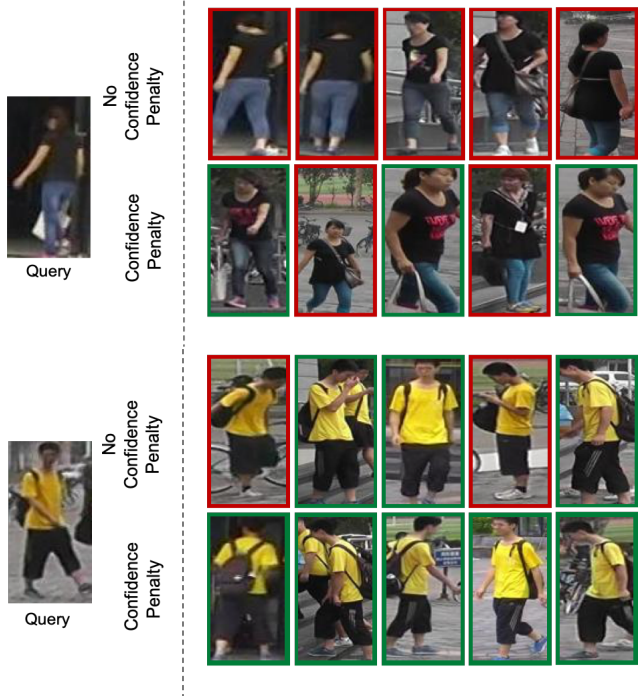


Figure 4: Qualitative comparison of using confidence penalty on unseen test samples. The gallery images are ranked according to L2 distance (top-5, left to right). Red frame indicates wrong ID while green frame indicates correct ID compared to the query. Best viewed in color.

individuals that the baseline tends to misidentify. These are ideal examples for why confidence penalty drastically improved person re-id compared to less significant improvements in other computer vision tasks.

7. Conclusions

We emphasize an intrinsic characteristic of person re-identification that poses a problem to the network being trained. The classes that person re-id tries to separate are not as easy as separating cats and dogs. Different people with different identities can have very similar appearances. We have demonstrated that three methods, that reduce a model’s confidence, are able to deal with this problem while achieving state-of-the-art results. Confidence penalty proved to be the best performant and most lightweight amongst the different methods. In addition, it is interesting to note that VIB is able to achieve similar results while using smaller representations. Both label smoothing and confidence penalty use a representation of 2048 while VIB uses a representation of size 1024. These three methods can be leveraged to improve the performance of previous re-id methods as well. It remains an exciting future work to study their effect on other image retrieval and clustering tasks.

References

- [1] E. Ahmed, M. Jones, and T. K. Marks. An improved deep learning architecture for person re-identification. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2015. 2, 3
- [2] A. A. Alemi, I. Fischer, J. V. Dillon, and K. Murphy. Deep variational information bottleneck. *CoRR*, abs/1612.00410, 2016. 2, 3, 5
- [3] X. Chang, T. M. Hospedales, and T. Xiang. Multi-level factorisation net for person re-identification. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2018. 2, 7
- [4] D. Chen, D. Xu, H. Li, N. Sebe, and X. Wang. Group consistent similarity learning via deep crf for person re-identification. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2018. 2, 7
- [5] W. Chen, X. Chen, J. Zhang, and K. Huang. Beyond triplet loss: A deep quadruplet network for person re-identification. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, July 2017. 3
- [6] D. Cheng, Y. Gong, S. Zhou, J. Wang, and N. Zheng. Person re-identification by multi-channel parts-based cnn with improved triplet loss function. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2016. 2
- [7] P. F. Felzenszwalb, R. B. Girshick, D. McAllester, and D. Ramanan. Object detection with discriminatively trained part-based models. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 32(9):1627–1645, Sept 2010. 5
- [8] J. Garca, N. Martinel, C. Micheloni, and A. Gardel. Person re-identification ranking optimisation by discriminant context information analysis. In *2015 IEEE International Conference on Computer Vision (ICCV)*, pages 1305–1313, Dec 2015. 2
- [9] R. Hadsell, S. Chopra, and Y. LeCun. Dimensionality reduction by learning an invariant mapping. In *2006 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR'06)*, volume 2, pages 1735–1742, June 2006. 3
- [10] H. Huang, D. Li, Z. Zhang, X. Chen, and K. Huang. Adversarially occluded samples for person re-identification. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2018. 7
- [11] M. M. Kalayeh, E. Basaran, M. Gkmen, M. E. Kamasak, and M. Shah. Human semantic parsing for person re-identification. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2018. 2, 7
- [12] D. P. Kingma and M. Welling. Auto-encoding variational bayes. *CoRR*, abs/1312.6114, 2014. 5
- [13] Q. Leng, R. Hu, C. Liang, Y. Wang, and J. Chen. Person re-identification with content and context re-ranking. *Multimedia Tools Appl.*, 74(17):6989–7014, Sept. 2015. 2
- [14] D. Li, X. Chen, Z. Zhang, and K. Huang. Learning deep context-aware features over body and latent parts for person re-identification. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, July 2017. 2
- [15] W. Li, R. Zhao, T. Xiao, and X. Wang. Deepreid: Deep filter pairing neural network for person re-identification. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2014. 2
- [16] W. Li, X. Zhu, and S. Gong. Harmonious attention network for person re-identification. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2018. 2, 7
- [17] J. Liu, B. Ni, Y. Yan, P. Zhou, S. Cheng, and J. Hu. Pose transferrable person re-identification. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2018. 3, 7
- [18] G. Pereyra, G. Tucker, J. Chorowski, L. Kaiser, and G. E. Hinton. Regularizing neural networks by penalizing confident output distributions. *CoRR*, abs/1701.06548, 2017. 2, 3, 4
- [19] R. Rama Varior, B. Shuai, J. Lu, D. Xu, and G. Wang. A siamese long short-term memory architecture for human re-identification. volume 9911, pages 135–153, 10 2016. 3
- [20] S. J. Reddi, S. Kale, and S. Kumar. On the convergence of adam and beyond. In *International Conference on Learning Representations*, 2018. 6
- [21] M. Saquib Sarfraz, A. Schumann, A. Eberle, and R. Stiefelhagen. A pose-sensitive embedding for person re-identification with expanded cross neighborhood re-ranking. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2018. 2, 7
- [22] F. Schroff, D. Kalenichenko, and J. Philbin. Facenet: A unified embedding for face recognition and clustering. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2015. 6
- [23] Y. Shen, H. Li, S. Yi, D. Chen, and X. Wang. Person re-identification with deep similarity-guided graph neural network. In *The European Conference on Computer Vision (ECCV)*, September 2018. 2, 7
- [24] Y. Shen, T. Xiao, H. Li, S. Yi, and X. Wang. End-to-end deep kronecker-product matching for person re-identification. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2018. 2, 7
- [25] J. Si, H. Zhang, C.-G. Li, J. Kuen, X. Kong, A. C. Kot, and G. Wang. Dual attention matching network for context-aware feature sequence based person re-identification. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2018. 7
- [26] C. Song, Y. Huang, W. Ouyang, and L. Wang. Mask-guided contrastive attention model for person re-identification. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2018. 2
- [27] Y. Suh, J. Wang, S. Tang, T. Mei, and K. Mu Lee. Part-aligned bilinear representations for person re-identification. In *The European Conference on Computer Vision (ECCV)*, September 2018. 7
- [28] C. Szegedy, W. Liu, Y. Jia, P. Sermanet, S. Reed, D. Anguelov, D. Erhan, V. Vanhoucke, and A. Rabinovich. Going deeper with convolutions. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2015. 2

- [29] C. Szegedy, V. Vanhoucke, S. Ioffe, J. Shlens, and Z. Wojna. Rethinking the inception architecture for computer vision. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2016. 3, 4
- [30] N. Tishby, F. C. Pereira, and W. Bialek. The information bottleneck method. pages 368–377, 1999. 4
- [31] R. R. Variator, B. Shuai, J. Lu, D. Xu, and G. Wang. A siamese long short-term memory architecture for human re-identification. In *Computer Vision – ECCV 2016*, pages 135–153. Springer International Publishing, 2016. 2
- [32] C. Wang, Q. Zhang, C. Huang, W. Liu, and X. Wang. Mancs: A multi-task attentional network with curriculum sampling for person re-identification. In *The European Conference on Computer Vision (ECCV)*, September 2018. 2, 3, 7
- [33] Y. Wang, L. Wang, Y. You, X. Zou, V. Chen, S. Li, G. Huang, B. Hariharan, and K. Q. Weinberger. Resource aware person re-identification across multiple resolutions. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2018. 2, 3, 7
- [34] L. Wei, S. Zhang, W. Gao, and Q. Tian. Person transfer gan to bridge domain gap for person re-identification. In *Computer Vision and Pattern Recognition, IEEE Conference on*, 2018. 5, 8
- [35] S. Wu, Y.-C. Chen, X. Li, A.-C. Wu, J.-J. You, and W.-S. Zheng. An enhanced deep feature representation for person re-identification. In *2016 IEEE Winter Conference on Applications of Computer Vision (WACV)*, pages 1–8. IEEE, 2016. 1
- [36] T. Xiao, H. Li, W. Ouyang, and X. Wang. Learning deep feature representations with domain guided dropout for person re-identification. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2016. 1
- [37] J. Xu, R. Zhao, F. Zhu, H. Wang, and W. Ouyang. Attention-aware compositional network for person re-identification. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2018. 2
- [38] H. Yao, S. Zhang, Y. Zhang, J. Li, and Q. Tian. Deep representation learning with part loss for person re-identification. *CoRR*, abs/1707.00798, 2017. 2
- [39] M. Ye, C. Liang, Z. Wang, Q. Leng, and J. Chen. Ranking optimization for person re-identification via similarity and dissimilarity. In *Proceedings of the 23rd ACM International Conference on Multimedia*, MM ’15, pages 1239–1242, New York, NY, USA, 2015. ACM. 2
- [40] M. Ye, C. Liang, Y. Yu, Z. Wang, Q. Leng, C. Xiao, J. Chen, and R. Hu. Person reidentification via ranking aggregation of similarity pulling and dissimilarity pushing. *IEEE Transactions on Multimedia*, 18(12):2553–2566, Dec 2016. 2
- [41] D. Yi, Z. Lei, S. Liao, and S. Z. Li. Deep metric learning for person re-identification. In *2014 22nd International Conference on Pattern Recognition*, pages 34–39, Aug 2014. 2
- [42] R. Yu, Z. Dou, S. Bai, Z. Zhang, Y. Xu, and X. Bai. Hard-aware point-to-set deep metric for person re-identification. In *The European Conference on Computer Vision (ECCV)*, September 2018. 3, 7
- [43] L. Zhao, X. Li, Y. Zhuang, and J. Wang. Deeply-learned part-aligned representations for person re-identification. In *The IEEE International Conference on Computer Vision (ICCV)*, Oct 2017. 2
- [44] L. Zheng, L. Shen, L. Tian, S. Wang, J. Wang, and Q. Tian. Scalable person re-identification: A benchmark. In *Computer Vision, IEEE International Conference on*, 2015. 5
- [45] Z. Zheng, L. Zheng, and Y. Yang. Unlabeled samples generated by gan improve the person re-identification baseline in vitro. In *Proceedings of the IEEE International Conference on Computer Vision*, 2017. 5
- [46] Z. Zhong, L. Zheng, D. Cao, and S. Li. Re-ranking person re-identification with k-reciprocal encoding. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, July 2017. 2, 7
- [47] Z. Zhong, L. Zheng, G. Kang, S. Li, and Y. Yang. Random erasing data augmentation. *CoRR*, abs/1708.04896, 2017. 6
- [48] Z. Zhong, L. Zheng, Z. Zheng, S. Li, and Y. Yang. Camera style adaptation for person re-identification. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2018. 7