Video Person Re-ID: Fantastic Techniques and Where to Find Them (Student Abstract)

Priyank Pathak,^{1*} Amir Erfan Eshratifar,^{2*} Michael Gormish^{3*}

¹Department of Computer Science, New York University, New York, NY 10012, USA

²University of Southern California, Los Angeles, CA 90089, USA

³Clarifai, San Francisco, CA 94105, USA

ppriyank@nyu.edu, eshratif@usc.edu, michael.gormish@clarifai.com

Abstract

The ability to identify the same person from multiple camera views without the explicit use of facial recognition is receiving commercial and academic interest. The current status-quo solutions are based on attention neural models. In this paper, we propose Attention and CL loss, which is a hybrid of center and Online Soft Mining (OSM) loss added to the attention loss on top of a temporal attention-based neural network. The proposed loss function applied with bag-of-tricks for training surpasses the state of the art on the common person Re-ID datasets, MARS and PRID 2011. Our source code is publicly available on github¹.

Introduction

Person Re-IDentification (Re-ID) aims to recognize the same individual in different pictures or image sequences caught by distinct cameras that may or may not be temporally aligned (Fig 1). The goal of video-based person Re-ID is to find the same person in a set of gallery videos from a query video. One industrial use case is surveillance for security purposes. In contrast with image-based person Re-id, which can be consequently affected by several factors such as blurriness, lighting, pose, viewpoint and occlusion, video-based person Re-ID is more robust to these pitfalls as multiple frames distributed across the video are used.

Methodology

Baseline (Base Temporal Attention): We use Gao and Nevatia’s (2018) temporal attention model as our base model, where a pre-trained ResNet-50 on ImageNet creates features for each frame of a video clip, and an attention model computes a weighted sum of the features across frames.

Bag-of-Tricks: Luo et al. (2019) proposed a series of tricks to enhance the performance of a ResNet model in image-based person Re-ID, which includes reducing the stride of

Figure 1: Data Samples are taken from MARS (Zheng et al. 2016). Same individual captured in 3 input video clips

the last layer (richer feature space), using warm-up learning rate, random erasing of patches within frames (simulating occlusion), label smoothing, center loss in addition to triplet loss, cosine-metric based triplet loss (angular distance proven better than L-2 distance) and lastly, batch normalization before the classification layer. Our experiments reveal that batch normalized feature provides a more robust model compared to non-normalized features.

Attention and CL loss: Wang et al. (2018) proposed Online Soft Mining and Class-Aware Attention (OSM loss) as an alternative to triplet loss for training Re-ID tasks, a modified contrastive loss with attention to remove noisy frames. We propose **CL Centers** OSM loss, which uses the center vectors from center loss as the class label vector representations, for cropping out noisy frames as they have greater variance compared to the originally proposed classifier weights. In addition, we penalize the model for giving high attention scores to frames where we have randomly deleted a patch. Such randomly erased frames are labeled as 1 otherwise 0.

$$\text{Attention loss} = \frac{1}{N} \sum_{i=1}^N \text{label}(i) * \text{Attention}_{\text{score}}(i) \quad (1)$$

where N is the number of total frames.

*Work done at Clarifai Inc.

Copyright © 2020, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

¹<https://github.com/ppriyank/Video-Person-Re-ID-Fantastic-Techniques-and-Where-to-Find-Them>

Figure 2: The proposed model architecture. $\otimes$ indicates pairwise multiplication and $\oplus$ indicate summation.

Model	mAP (re-rank)	CMC-1 (re-rank)	CMC-5 (re-rank)	CMC-20 (re-rank)
SOTA w/o re-ranking (Fu et al., (2018))	81.2 (-)	86.2(-)	95.7(-)	- (-)
SOTA with re-ranking (Fu et al., (2018))	80.8 (87.7)	86.3(87.2)	95.7(96.2)	98.1 (98.6)
Baseline (Gao and Nevatia, (2018))	76.7 (84.5)	83.3 (85.0)	93.8 (94.7)	97.4 (97.7)
Baseline + Bag-of-Tricks (B-BOT)*	81.3 (88.4)	87.1 (87.6)	95.9 (96.0)	98.2 (98.4)
B-BOT + OSM loss (B-BOT + OSM)*	82.4 (88.1)	87.9 (87.6)	96.0 (95.7)	98.0 (98.5)
(Proposed) B-BOT + OSM + CL Centers*	81.2 (88.5)	86.3 (88.0)	95.6 (96.1)	98.2 (98.5)
(Proposed) B-BOT + Attention and CL loss*	82.9 (87.8)	88.6 (88.0)	96.2 (95.4)	98.0(98.3)

Table 1: MARS Dataset Performance. '-' indicates the results were not reported. '*' refers to hyperparameter optimized.

Model	CMC-1	CMC-5	CMC-20
SOTA (Zeng, Tian, and Wu, (2018))	96.1	99.5	-
(Proposed) B-BOT + Attn-CL loss*	93.3	98.9	100.0
(Proposed) B-BOT + Attn-CL loss (pre-trained on MARS dataset)*	96.6	100	100.0

Table 2: PRID 2011 Dataset Performance. '-' indicates the results were not reported. '*' refers to hyperparameter optimized.

The Attention loss combined with OSM loss and CL Centers is denoted as **Attention and CL loss**.

Hyperparameter optimization: We also applied Facebook's hyperparameter optimization tool² to do hyperparameter search.

Datasets: We focus on the MARS and PRID datasets containing 1251 and 178 identities, respectively which are equally split among training and testing sets.

Evaluation and Results

In our experiments, we use four frames of the video selected randomly, $N = 4$. Figure 2 shows the proposed model architecture. Table 1 and Table 2 show a comparison of our model to the state-of-the-art results on MARS and PRID 2011 datasets.

Conclusion and Future Work

In this paper, we mixed and improved existing techniques to surpass the state-of-the-art accuracy on MARS and PRID 2011 datasets. We plan to evaluate our work on other datasets and other similar tasks like facial re-identification in the future.

²<https://github.com/facebook/Ax>

References

- Fu, Y.; Wang, X.; Wei, Y.; and Huang, T. 2018. STA: spatial-temporal attention for large-scale video-based person re-identification. *CoRR* abs/1811.04129.
- Gao, J., and Nevatia, R. 2018. Revisiting temporal modeling for video-based person reid. *CoRR* abs/1805.02104.
- Luo, H.; Gu, Y.; Liao, X.; Lai, S.; and Jiang, W. 2019. Bag of tricks and A strong baseline for deep person re-identification. *CoRR* abs/1903.07071.
- Wang, X.; Hua, Y.; Kodirov, E.; Hu, G.; and Robertson, N. M. 2018. Deep metric learning by online soft mining and class-aware attention. *CoRR* abs/1811.01459.
- Zeng, M.; Tian, C.; and Wu, Z. 2018. Person re-identification with hierarchical deep learning feature and efficient xqda metric. In *Proceedings of the 26th ACM International Conference on Multimedia*, MM '18, 1838–1846. New York, NY, USA: ACM.
- Zheng, L.; Bie, Z.; Sun, Y.; Wang, J.; Su, C.; Wang, S.; and Tian, Q. 2016. Mars: A video benchmark for large-scale person re-identification. In *ECCV*.