Communicated by Dr. Qingshan Liu

Accepted Manuscript

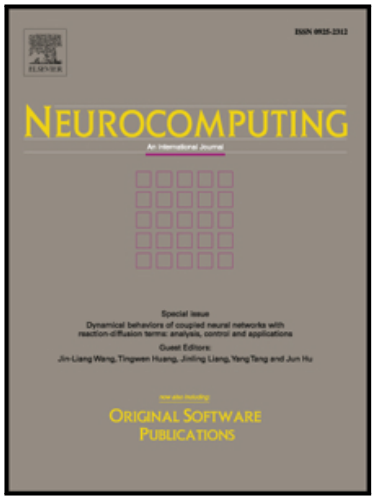

Improving Face Recognition with Domain Adaptation

Ge Wen, Huaguan Chen, Deng Cai, Xiaofei He

PII: S0925-2312(18)30112-7
DOI: 10.1016/j.neucom.2018.01.079
Reference: NEUCOM 19281

To appear in: *Neurocomputing*

Received date: 19 October 2017
Revised date: 8 January 2018
Accepted date: 28 January 2018

Please cite this article as: Ge Wen, Huaguan Chen, Deng Cai, Xiaofei He, Improving Face Recognition with Domain Adaptation, *Neurocomputing* (2018), doi: 10.1016/j.neucom.2018.01.079

Improving Face Recognition with Domain Adaptation

Ge Wen, Huaguan Chen, Deng Cai*, Xiaofei He

The State Key Lab of CAD&CG, Zhejiang University, No.388 Yu Hang Tang Road, Hangzhou 310058, China

Abstract

Nearly all recent face recognition algorithms have been evaluated on the Labeled Faces in the Wild (LFW) dataset and many of them achieved over 99% accuracy. However, the performance is still not enough for real-world applications. One problem is the data bias. The faces in LFW and other web-collected datasets come from celebrities. They are quite different from the faces of a normal person captured in the daily life. In other words, they are different in the face distribution. Replacing the training data with the right distribution is a simple solution. However, the photos of common people are much harder to collect because of the privacy concerns. So it is useful to develop a method that transfers the knowledge in the data of different face distribution to help improving the final performance. In this paper, we crawl a large face dataset whose distribution is different from LFW and show the improvement of LFW accuracy with a simple domain adaptation technique. To the best of our knowledge, it is the first time that domain adaptation is applied in the unconstrained face recognition problem with million scale dataset. Besides, we incorporate face verification threshold into FaceNet triplet loss function explicitly. Finally, we achieve 99.33% on the LFW benchmark with only single CNN model and similar performance even without face alignment.

Keywords: Face recognition, domain adaptation, face verification loss

Introduction

Face recognition is the problem of identifying a specific individual, rather than merely detecting the presence of a human face. It is widely used in public security, finance security, commercial domain and so on. Due to its wide applications, face recognition has become a core problem and one of the most popular research topics in computer vision. It includes two different but related tasks, face verification (are these two pictures the same person) and face identification (who is this person). Face verification can be extended to solve face identification task by repeating one-vs-one comparison. Nearly all recent methods have been evaluated on the Labeled Faces in the Wild (LFW) dataset [8]. In this paper, we focus on the face verification task and report the performance on the LFW benchmark as well.

The recent face recognition methods are based on convolutional neural network [13], and have made a great progress, even beating human beings on the LFW benchmark. But its performance is still not enough for real-world applications. One problem is the data bias[32]. The faces in LFW and other web-collected datasets come from celebrities. They are quite different from the faces of a normal person captured in the daily life. In other words, they are different in the face distribution. Replacing the training data with the right distribution is a simple solution. But the photos of common people are much harder to collect because of the privacy concerns. Besides, a generic recognition system is required to be transferred to a domain-specific application for performance. Both can be formulated as Domain Adaptation[19], which transfers the knowledge in the source domain to the target domain. In this paper, we crawl a large face dataset called TaoMM whose distribution is different from LFW and show the improvement of LFW accuracy with a simple domain adaptation technique.

In the testing phase of face verification, the distance between face pair is compared with a pre-computed threshold θ. If $dis < \theta$, the face pair is regarded as from the same person, otherwise from different person. There is a similar threshold θ in open-set face identification. Most face recognition methods don't consider the threshold in their optimization process explicitly. So there exists an optimization gap in their methods. By incorporating the threshold into FaceNet [23] triplet loss function explicitly, we reduce the LFW error rate by 26.9%. DDML[7] use a similar idea, but our final formulation is a triplet loss in an end-to-end framework.

Data augmentation is a very common preprocessing step for CNN based method[11], as a CNN model contains millions of parameters and is prone to overfitting. Most face recognition methods [29, 26, 25, 27, 4, 30] align face in both training and testing phase. It seems contradictory to apply data augmentation after face alignment. In this paper, we replace face alignment in training phase with aggressive data augmentation. Surprisingly, similar accuracy is achieved on LFW benchmark with or without face alignment during testing, which is different from prior results[18, 23].

*Corresponding author

Email addresses: zjuwenge@gmail.com (Ge Wen), chenhuaguanzju@gmail.com (Huaguan Chen), dengcai@cad.zju.edu.cn (Deng Cai), xiaofeihe@cad.zju.edu.cn (Xiaofei He)