

Міністерство освіти і науки України
НАЦІОНАЛЬНИЙ УНІВЕРСИТЕТ «КИЄВО-МОГИЛЯНСЬКА АКАДЕМІЯ»
Кафедра інформатики факультету інформатики



Курсова робота
за спеціальністю „Інженерія програмного забезпечення”

**Face recognition in the video stream. Self-attention neural
aggregation network**

Керівник курсової роботи
доц. к ф. -м н, Крюкова Г. В.

(підпис)

“ ____ ” _____ 2019 р.

Виконав
студент 4 курсу
факультету інформатики
Проценко І. В.

Work Schedule

Number	Coursework writing stages	Date
1	Research on metric learning approaches and neural aggregation architectures	01.12.2019
2	Dataset searching, annotation processing, facial descriptors extraction for single identity aggregation protocol.	31.11.2019
3	Architecture creation, training and validation	25.12.2019
4	Multi-identity approach adaption. Synthetic dataset creation.	17.02.2020
5	Completed coursework draft	01.04.2029
6	Completed final coursework draft	15.04.2020

Table of contents

Table of contents	3
Abstract	4
Introduction	4
Related Work	6
Image Face Recognition	6
Video Face Recognition	6
Aggregation networks	8
Single-identity aggregation	9
Multi-identity aggregation	11
Experiments	13
Datasets and protocols	13
Preprocessing	14
Training details	14
Single identity training	14
Multi-identity training	15
Results for single-identity aggregator	16
UMD+GIFs validation set	16
IJB-C	17
Results for multi-identity	19
UMD+GIFs validation set	19
IJB-C	19
Conclusions	21
References	21

Abstract

The models based on self-attention mechanisms have been successful in analyzing temporal data and have been widely used in the natural language domain. A new model architecture is being proposed for video face representation and recognition based on the self-attention mechanism. Moreover, given approach could be used for video with single and multiple identities. Notably, no one explored the aggregation approaches that consider the video with multiple identities. The proposed approach utilizes existing models to get the face representation for each video frame, e.g., ArcFace and MobileFaceNet, and the aggregation module produces the aggregated face representation vector for video by taking into consideration the order of frames and their quality scores. Empirical results are demonstrated on a public dataset for video face recognition called IJB-C to indicate that the self-attention aggregation network (SAAN) outperforms naive average pooling. Moreover, a new multi-identity video dataset based on the publicly available UMDFaces dataset and collected GIFs from Giphy is being proposed. It is shown that SAAN is capable of producing a compact face representation for both single and multiple identities in a video. The source code is attached in the archive.

Introduction

Video face recognition has received increasing interest from the community in recent years [11, 12, 24, 28, 30, 36, 38, 46, 42], partially due to the growing number of video content. Compared to an image, a video sequence contains more information about the subject's faces, such as varying poses, expressions, motion, and illumination. The key challenge for video-based face recognition is how to effectively combine facial information available across different video frames to get an appropriate representation of the face in the video.

Deep Convolutional Neural Networks (DCNNs) trained on a large dataset have shown the ability to generate compact and discriminative face representations for images that are

robust to pose variations, image quality, blur, and occlusions [31, 42, 45, 46]. The naive approach is to represent video as a set of frames and use face features extracted by deep neural network for each frame [42, 46]. Then the subject's face in the video, called face track, is represented as an unordered set of vectors that allows maintaining the information across all frames. However, this is not computationally efficient for comparing two face tracks. It requires a fusion of matching results and comparison across all pairs of vectors between two face templates. Except for $O(n^2)$ complexity per match operation, face track requires $O(n)$ space per video face, where n - the average number of video frames. That is why most methods aggregate feature vectors into fix-size feature representation for each face track [33, 41, 40, 44, 54, 55]. It allows constant-time matching computation instead of evaluating all powers.

In this coursework, the aggregation model based on self-attention is proposed, which can be applied for videos that contain single or multiple identities. Also, the synthetic dataset is being created for multi-identity video face recognition.

Related Work

Image Face Recognition

Face recognition is an actively studied domain with significant achievements in identification and verification tasks [48, 50, 58, 31] and a great part of that success is due to deep convolutional neural networks. Most approaches are focused on learning the embedding vector for face representation. The following works focus on exploring different loss functions to improve the feature representation. Both contrastive [9, 21] and triplet [23, 51] losses are usually used to increase the Euclidean margin for better feature embedding.

The center loss proposed in [53] learns centroids for features of each identity to reduce intra-class variance. A large margin softmax (L-softmax) [32] adds angular constraints to each identity to improve feature discrimination, and angular softmax (A-softmax) [31] adds weights normalization for L-softmax. The ArcFace uses additive angular margin loss [13].

Video Face Recognition

Existing methods could be classified into the following categories: ones that exploit temporal dynamics and ones that treat video as an orderless set of images. The first group of methods heavily relies on RNNs to account for the temporal dependencies in frame sequences. For example, RNN was employed for head pose estimation [19], facial expression recognition [18], and emotion recognition [15, 59]. Many previous methods have considered the representation of the set of face images as probabilistic distribution [2, 43], n-order statistics [34], affine hulls [57, 25, 5], SPD matrices [26], manifolds [27, 22, 52], etc. Then, the recognition is performed via similarity or distance measures. Other methods aim to train a supervised classifier on each image set or video to obtain correspondent representation [37].

These methods work well under constrained settings but usually struggle to handle the challenging unconstrained scenarios with significant appearance variations. The aggregation based models aim to fuse a set of deep feature vectors into a single vector. It was shown in [4] that aggregation of multiple face images increases the recognition performance of person identification. Compared to simple average pooling [7, 14, 6, 10,

[13], the recent works show promising results by incorporating the visual quality information on instance level via detection score [39] or predicted quality scores [55, 56]. A Component-wise Feature Aggregation Network [17] aggregates the feature vectors in each component separately by considering the prediction of their quality. The redundancy issue in the video frames is tackled in [16]. However, none of these approaches consider the multi-identity video recognition setup.

Aggregation networks

In the following sections, we'll review proposed single and multi-identity aggregation architectures. To our knowledge, the aggregation of multiple identities has not been addressed so far. We consider a multi-head self-attention mechanism that has been successfully used to encode/decode sequence representation as in [47], achieving superior results and better parallelization compared to recurrent encoder/decoder framework. It is used to achieve reweighting according to the context of the features.

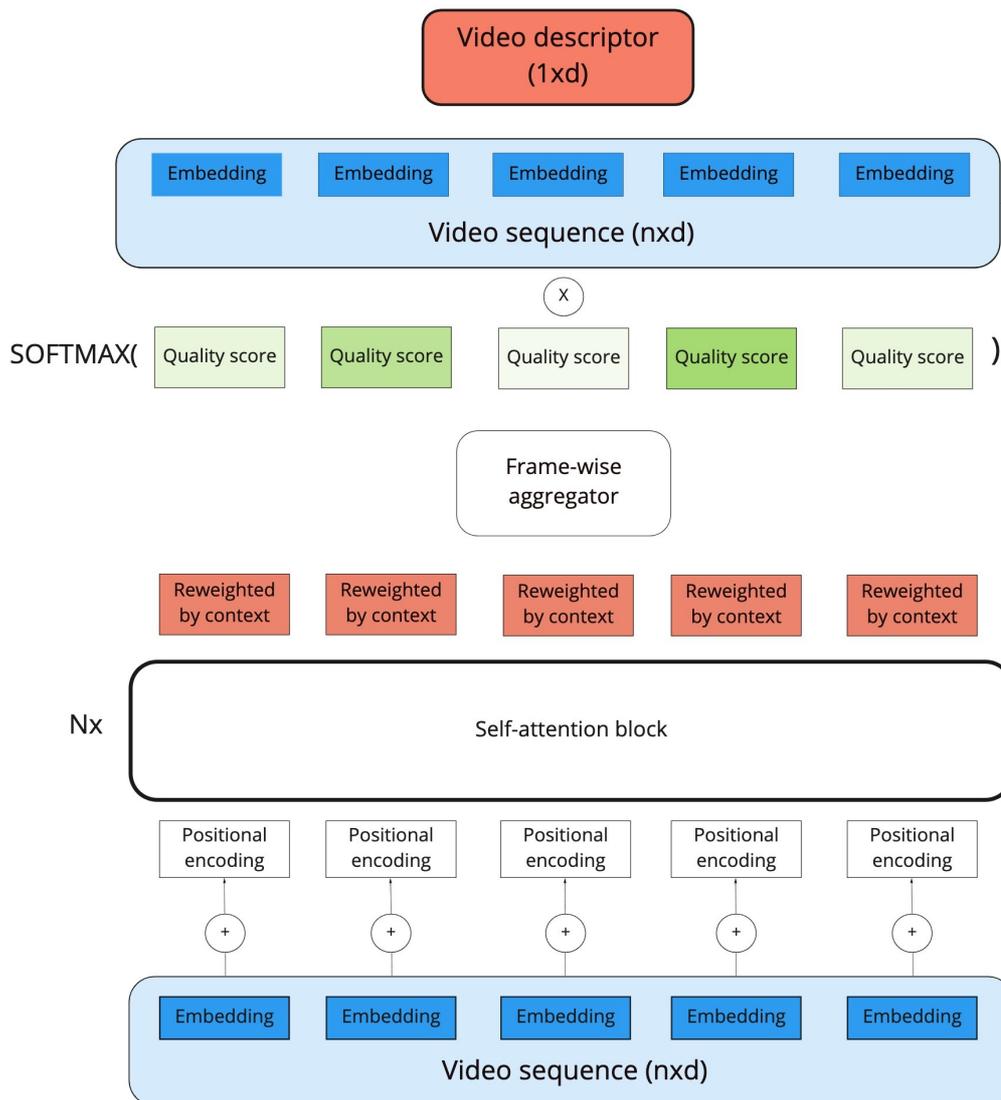


Fig1. Example of video sequence aggregation network using self-attention. n is the number of faces in a face track, d is a dimensionality of the embedding, N - is the number of consecutive self-attention blocks of the encoder. Figure is read in a bottom-up fashion.

To get a compact representation for each image in a given set, we publicly available **Arc-Face** model is used pretrained on the **MS1MV2** dataset. It is considered to be state-of-the-art in a lot of public verification/identification benchmarks.

Additionally, there are experiments with a lightweight feature extractor, in particular, **MobileFaceNet**. The main purpose of that experiment is to investigate the aggregation results on features with a lower representational capacity.

Single-identity aggregation

Given a set of ordered frames $\{x_1, x_2, x_3, x_4, \dots, x_n\}$ (face track) and their corresponding features $\{f_1, f_2, f_3, f_4, \dots, f_n\}$, aggregation should be performed to produce a single vector representation r . To represent a face track from different perspectives and assign weight to a particular frame with respect to its context, we resort to multi-head self-attention mechanism

At first, K, V, Q projections are computed by learnable transformation matrices on $(F+P)$, where F is an extracted embedding matrix, and P is a positional encoding matrix. Later, reweighted embeddings computed as:

$$RHead = \text{Softmax}\left(\frac{QK^T}{\sqrt{d}}\right)V \quad (1)$$

$$RMultiHead = \text{Concat}(RHead_0, RHead_1, \dots, RHead_n)W_o$$

where d - is the dimensionality of a feature vector, **RHead** - reweighted representation received from single attention head, **RMultiHead** - final reweighted representation from

multiple heads and \mathbf{W}_o - output transformation matrix. Each reweighted projection is assigned an element-wise or component-wise score (as in [55] or [17]).

According to those scores, aggregation is done on the original features in the following way:

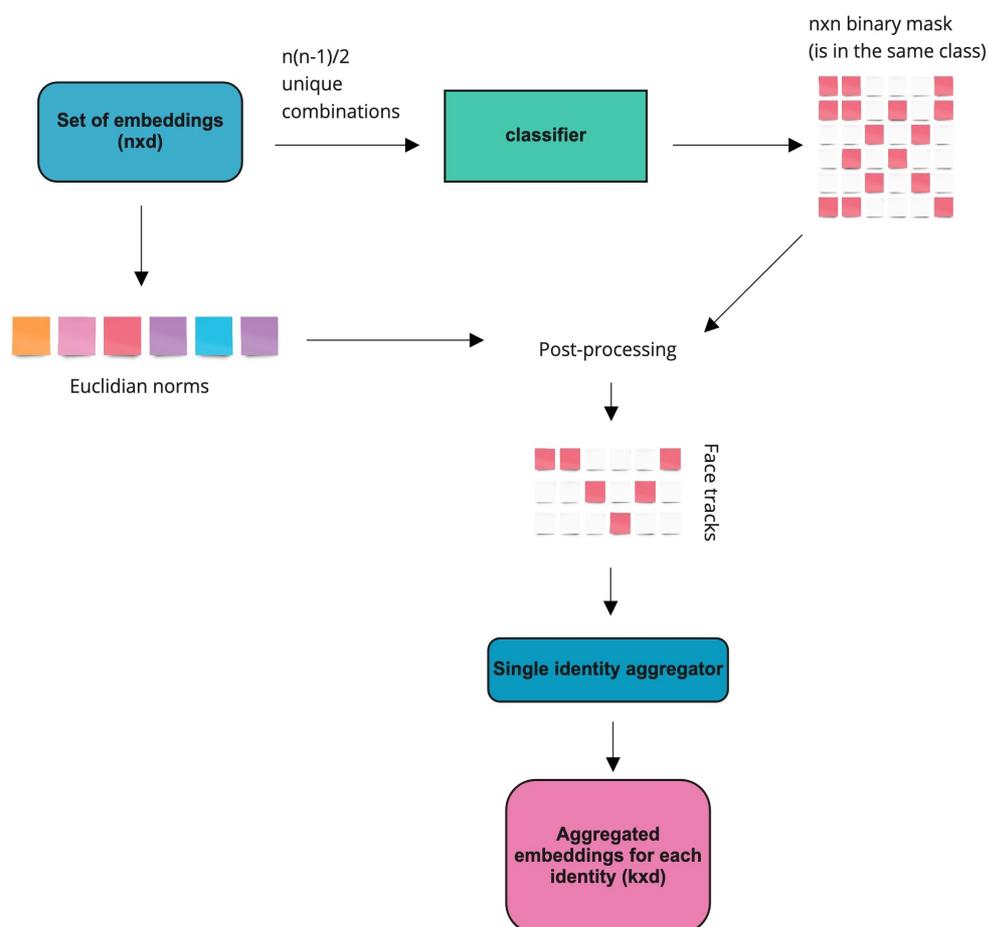
$$\mathbf{s} = RMultiHead \cdot \mathbf{W}_q \quad \mathbf{q} = Softmax(\mathbf{s})$$
$$\mathbf{r} = \sum_{i=1}^n f_i * \mathbf{q}_i$$

where \mathbf{s} is a vector of unnormalized quality scores, \mathbf{q} is a vector of softmax normalized scores, and \mathbf{W}_q is a learnable matrix to retrieve the quality score of given projection.

Multi-identity aggregation

The formulation of the task is similar to the previous section with the output being a set of vectors $[r_1, r_2, \dots, r_k]$ where k is the number of unique identities within a video sequence. Given a set of n features, multiple identities aggregation could be decomposed into two parts:

1. finding k face tracks within a set.
2. aggregating features within each face track.



Proposed multi-identity architecture. n - specifies the number of embeddings in a video, d - dimensionality of the embedding, k - number of unique identities within a video

By representing relation to the face tracks as a (n,n) binary matrix, our goal is to find a set of cliques. First, to create a binary mask, we need to classify all possible feature pairs within a video (that is $\binom{n}{2}$). After that, to remove any ambiguity and create a face track, we need to separate them into strongly connected components. To do that we implement a greedy post-processing. According to [39] norm of the embedding could be considered as a proxy to the quality of face embedding. Assignment of components is performed in the descending order specified by the Euclidian norm of the embeddings:

$norm_i = ||X_i||_2$, where \mathbf{X} is a matrix of extracted embeddings. All elements which lie in a class relation with a given example i are retrieved $\forall j \in \mathbf{R}(X_i, X_j)$, where \mathbf{R} is a class relation, assigned to the component and zeroed out in the initial binary matrix. The code snippet on Python is below.

```

1 import numpy as np
2
3 def postprocess_mask(p_mask, q_scores):
4     '''
5     p_mask :np.array(nxn) predicted mask
6     q_scores : np.array(nx1) l2 norms of input
7               embeddings
8     '''
9     p_mask_c = p_mask.copy()
10    f_mask = np.zeros_like(p_mask)
11    q_sorted = np.argsort(q_scores)
12
13    for q_ind in q_sorted[::-1]:
14        f_mask[q_ind] = p_mask_c[q_ind]
15        non_zero_els = np.argwhere(
16            f_mask[q_ind] != 0
17        )
18        p_mask_c[:, non_zero_els] = 0
19        p_mask_c[non_zero_els, :] = 0
20
21    return f_mask

```

Starting from the embedding with the highest norm, cliques are greedily assigned, until no elements are left unassigned. The resulting matrix is converted to the batch of zero-padded face tracks and then aggregated using a single identity aggregation network.

Experiments

Datasets and protocols

To train our aggregator, we use the **UMDFaces** [1] dataset. Additionally, we append a dataset with short GIFs of different identities parsed from publicly available API through **giphy.com**. Collected and preprocessed dataset will be made publicly available. For GIFs collection, we used their public API to get the most relevant GIFs for approximately 2300 celebrities. From each GIF, we sampled 5 frames evenly distributed in time. As a result, on average, we obtained about 150 GIFs for each identity, in total, almost 3M frames.

Sampled face tracks and overall training datasets statistics could be observed in the following figures:



Figure 3. Face tracks sampled from GIFs dataset.

Dataset	Frames	Identities	Sessions	Length
UMDFaces	3735476	3107	22075	161
GIFs	637939	2304	137739	4.6

Table 1. Training datasets descriptions. Length specified as an average number of frames per session. GIFs statistics correspond to the post-processed dataset.

IJB-C [35] dataset is used for benchmarking. IJB-C has 3,531 identities with 31.3K still images and 117.5K frames from 11, 779 videos.

Preprocessing

In the UMDFaces dataset, faces are detected using publicly available MTCNN model [58]. Using 5 facial keypoints, the similarity transform is applied to align faces. To match all detected faces with UMDFaces annotations, an IOU of 0.4 is chosen as a threshold. All other faces that were not matched are ignored. We use the same detector for GIFs dataset. It is important to note, that since GIF may contain multiple faces and not just the ones of desired celebrity, we used pretrained ArcFace model to get feature vectors for detected faces, computed mean vectors for each class and filtered out faces whose vectors do not lie within specified cosine distance threshold of 0.7 from the corresponding class center.

Training details

Single identity training

The aggregator is trained on the UMDFaces+GIFs dataset using an additive angular margin loss [49]. The radius of the hypersphere and additive margin are set to 16.0 and 0.35, respectively. The parameters are trained using a RAdam optimizer [29] with default values. For the encoder, the official implementation of Transformer on Tensorflow is used (<https://github.com/tensorflow/models/blob/master/official/transformer/model/transformer.py>), with a single modification (removing the trainable embedding block). The depth of the encoder is set to 4, the number of heads for the aggregator on top of ResNet and MobileNet is set to 8 and 4, respectively. Attention and ReLU dropout rates are specified as 0.3 and 0.4. All other hyperparameters are populated with the default values.

For validation, we employ an identity-based split. Due to the fact that identities in the validation are not present in the training set, we define a metric for the early stopping - intra-class proximity gain, defined as:

$$ICPG = E(\text{Intra}(\text{AVE}(X))) - E(\text{Intra}(\text{SA}(X)))$$

where X - is the set of templates for aggregation, SA and AVE are self-attention pooling and average pooling respectively, $\text{Intra}(Y)$ - is a set of distances between the aggregated elements of Y which belong to the same class. Specifically, each mini-batch includes 256 templates

that are randomly sampled from 128 identities, 32 images per template.

Multi-identity training

There are no multi-identity annotated video datasets that we are aware of, so we synthetically create one (combining UMDFaces and GIFs). We sample a random number of identities (from 2 to 64) and take a single session for each identity. These sessions are concatenated into a single one, which is then used for sampling 256 frames. These frames form a multi-identity video. The order of frames is preserved and used later in a single face track creation.

For the binary mask generation, there could be a trainable classifier, but we have found that simple thresholding based on the cosine distance results in higher recall and precision. To train the identity aggregator on top of the post-processed mask, we employ a teacher forcing with a scheduled sampling [3]. For the first 5000 iterations, cosine decay scheduler is used to specify the probability of mixing ground truth masks with the post-processed one. After that, only a predicted mask is used for the aggregation.

Results for single-identity aggregator

UMD+GIFs validation set

On the validation set, we report ICPG for the aggregators on top of different feature extractors as show in Table 2. Additionally, we attach visualization of attention extremums on the given set in Figure 4.

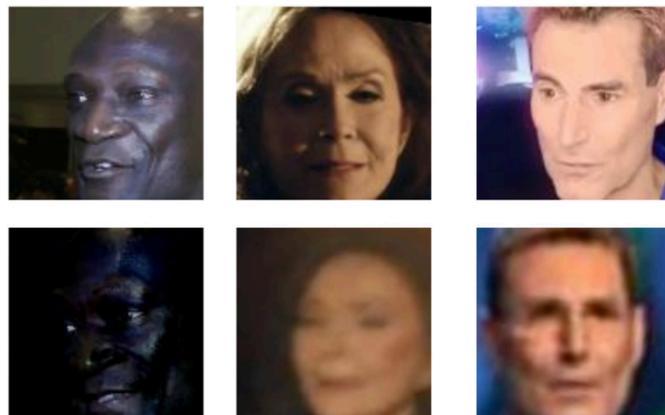


Figure 4. Attention extremums within the face tracks on the validation set. The first row represents the highest values, the second one - the lowest.

Aggregator	ICPG
SA pooling (ResNet)	$2.8 * 1e^{-3}$
SA pooling (MobileNet)	$5.5 * 1e^{-3}$

Table 2. ICPG on the validation set

It was found that there are some face tracks in the UMDFaces which contain erroneously assigned elements. Also, it is demonstrated in Figure 5. that by resorting to the general context, the self-attention mechanism helps to downweight such frames during final aggregation (despite their high visual quality).

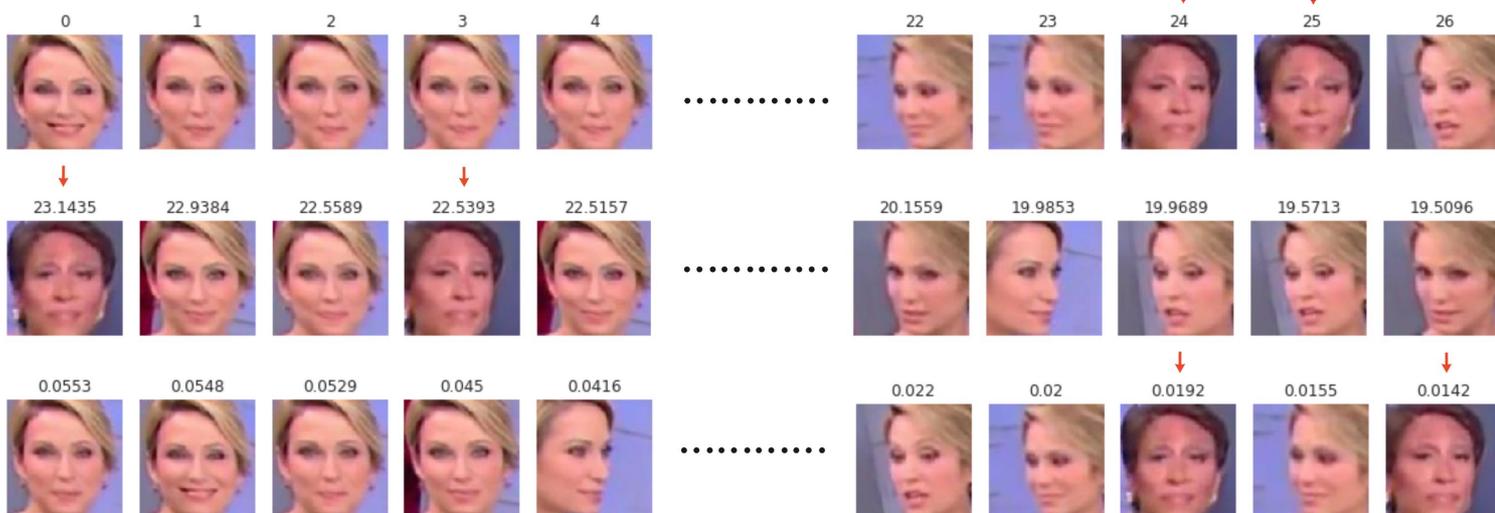


Figure 5. Attention distribution on the face track with erroneous elements. The first row - ordered frames of the given face track. The second row - frames ordered by the Euclidian norm of their corresponding embedding. The third row - frames ordered according to their self-attention quality score. The elements which do not belong to the given identity are emphasized with the red arrow.

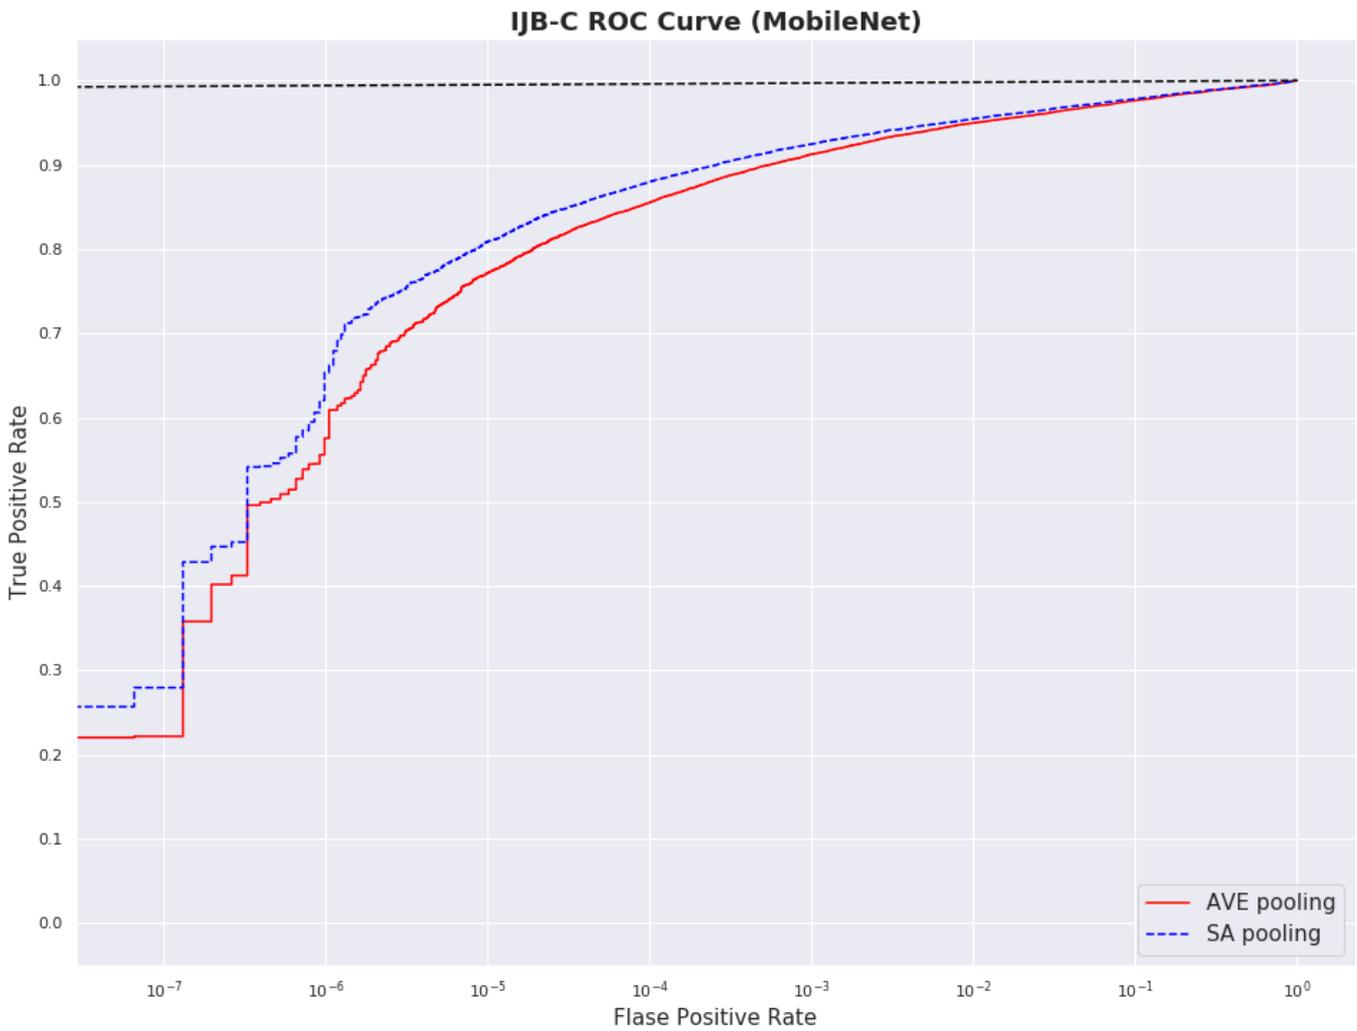
IJB-C

1:1 Verification on mixed set and 1:N end-to-end video probes are tested. AVE (average) and SA (self-attention) poolings are compared and results shown in the Table 3 and Figure 5.

The proposed approach results in a higher AUC, TAR, top-1 accuracy for both aggregators based on MobileNet and ResNet feature extractors. Notably, we observe that there is a more significant gap between aggregation and averaging based on the MobileNet feature extractor. It shows that the aggregation is more beneficial on the networks with the lower representational capacity.

Aggregator	Ver		Id
	TAR	AUC	Rank-1
AVE pooling (ResNet)	83.90	99.39	82.54
SA pooling (ResNet)	87.83	99.44	82.79
AVE pooling (MobileNet)	55.58	98.92	66.73
SA pooling (MobileNet)	62.04	99.01	67.48

Table 3. Verification TAR is reported under $FAR=1e-6$, identification rate is reported as rank-1 accuracy.



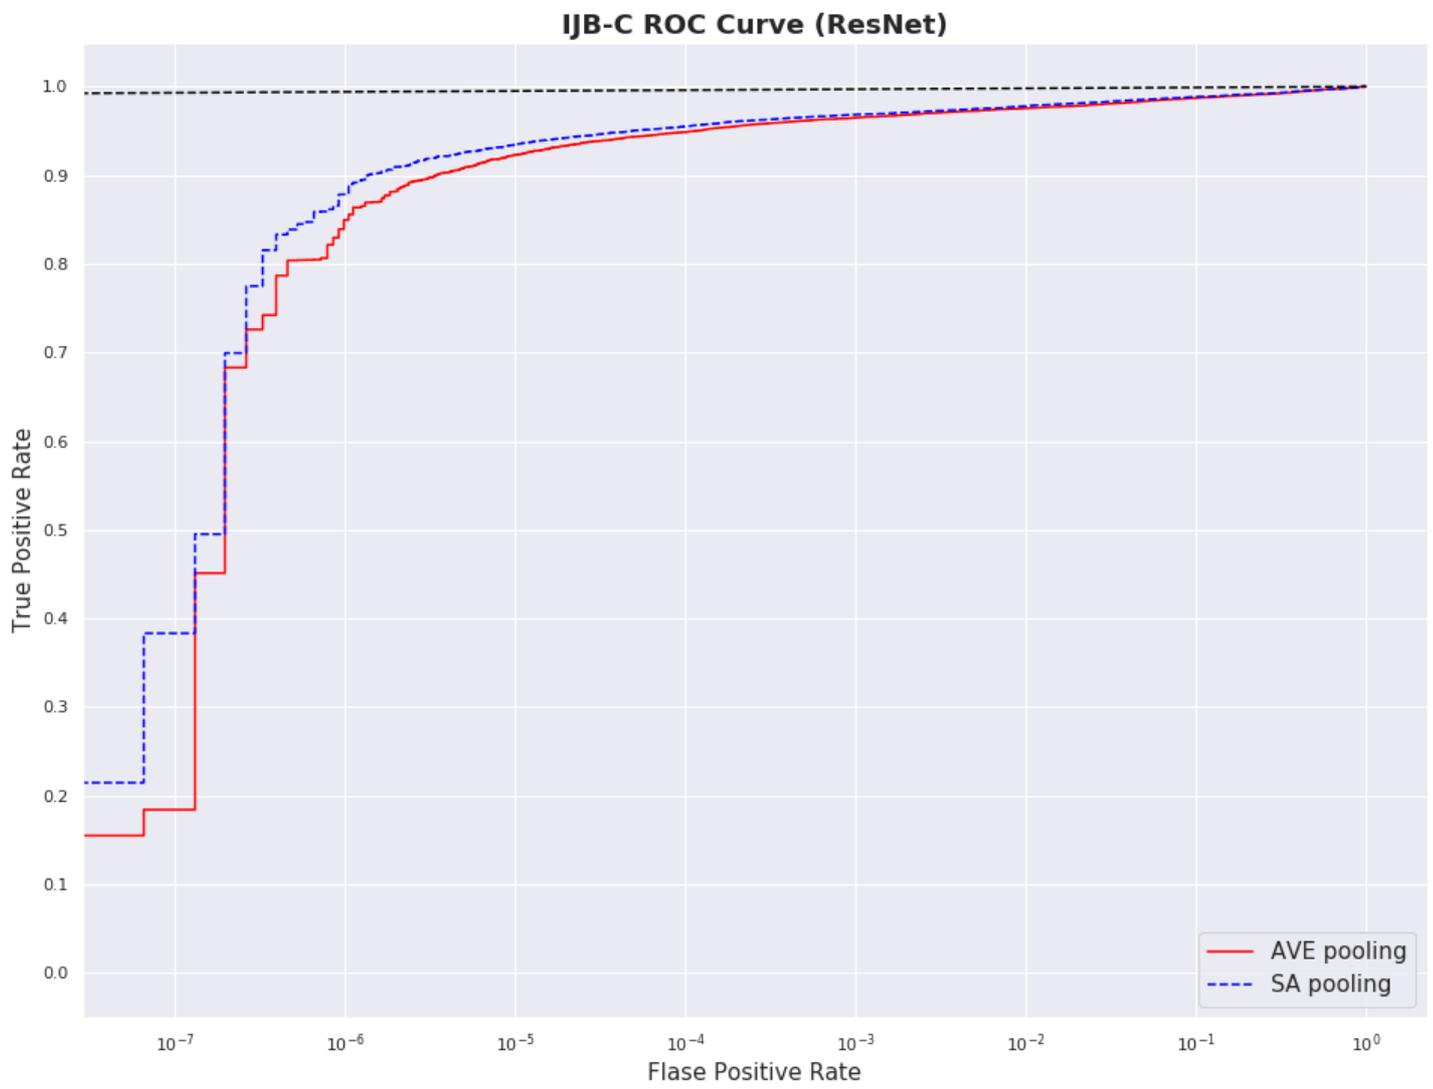


Figure 6. ROC curves of SA Pooling vs AVE Pooling on ResNet (top) and MobileNet (bottom) feature extractors.

Results for multi-identity

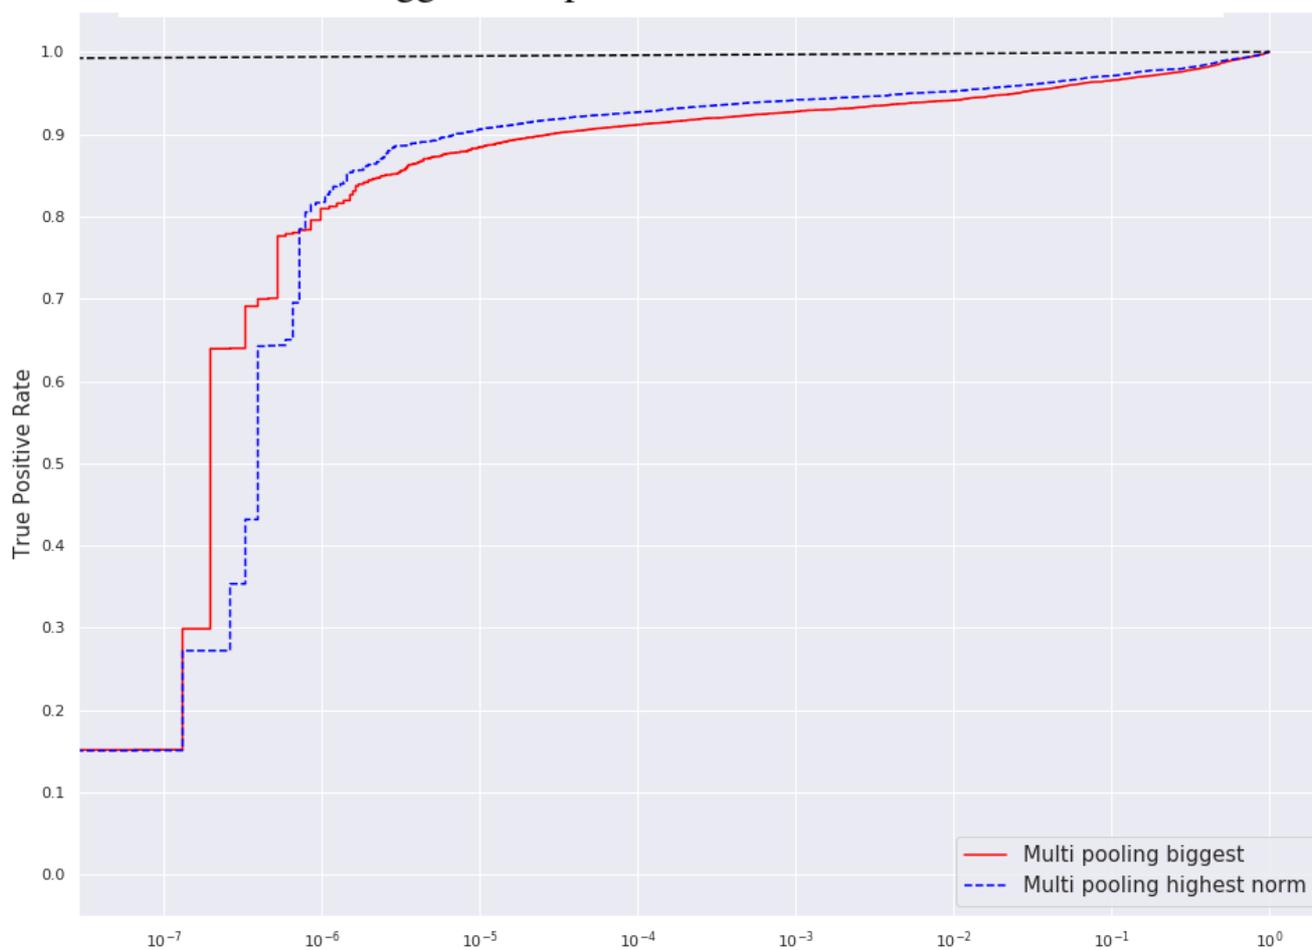
UMD+GIFs validation set

The number of regressed identities is calculated on the proposed multi-identity videos. Above mentioned binary mask producer results in MPE (mean percentage error) of **4.2%**. We presume that using different post-processing techniques for masks, such as transitive closure, could be used to lower that value.

IJB-C

Having only a single identity within a session, we introduce two component selection strategies (in order to select a single representation for reference in the session where multiple identities were detected). Component with largest amount of frames and component which contains the embedding of the highest quality are tested. Results are reported in Table 4 and ROC curve is displayed on Figure 7.

Figure 7. ROC curves of multi-identity aggregator with the highest norm and the biggest component selection.



Multi-identity aggregator	Ver		Id
	TAR	AUC	Rank-1
Highest norm component	81.67	98.58	82.54
Biggest component	79.55	98.30	82.02

Table 4. Verification TAR is reported under $FAR=1e-6$, identification rate is reported as rank-1 accuracy. Different component selection strategies are compared.

Conclusions

In this coursework, the self-attention aggregation network was proposed for learning face representation for any number of identities from a video stream. It is shown that SAAN outperforms average pooling in general for a single identity. Especially, the usage of SAAN network could be beneficial if vector representation quality degrades, e.g. when using lightweight embeddings from MobileNet. Moreover, the further investigation indicates that SAAN model is robust to erroneous face tracks. The future work will explore the different mask postprocessing approaches and ways to improve the aggregation model for multi-identity video streams.

References

1. Carlos Castillo, Rajeev Ranjan, Rama Chellappa, Ankan Bansal, Anirudh Nanduri. Umdfaces: An annotated face dataset for training deep networks. *arXiv preprint arXiv:1611.01484*, 2016. [3](#)
2. Ognjen Arandjelovic, Gregory Shakhnarovich, John Fisher, Roberto Cipolla, and Trevor Darrell. Face recognition with image sets using manifold density divergence. In *2005 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR'05)*, volume 1, pages 581–588. IEEE, 2005. [2](#)
3. Samy Bengio, Oriol Vinyals, Navdeep Jaitly, and Noam Shazeer. Scheduled sampling for sequence prediction with recurrent neural networks. In *Advances in Neural Information Processing Systems*, pages 1171–1179, 2015. [4](#)
4. Lacey Best-Rowden, Hu Han, Charles Otto, Brendan F Klare, and Anil K Jain. Unconstrained face recognition: Identifying a person of interest from a media collection. *IEEE Transactions on Information Forensics and Security*, 9(12):2144–2157, 2014. [2](#)
5. Hakan Cevikalp and Bill Triggs. Face recognition based on image sets. In *2010 IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, pages 2567–2573. IEEE, 2010. [2](#)
6. Jun-Cheng Chen, Rajeev Ranjan, Amit Kumar, Ching-Hui Chen, Vishal M Patel, and Rama Chellappa. An end-to-end system for unconstrained face verification with deep convolutional neural networks. In *Proceedings of the IEEE international conference on computer vision workshops*, pages 118–126, 2015. [2](#)
7. Jun-Cheng Chen, Rajeev Ranjan, Swami Sankaranarayanan, Amit Kumar, Ching-Hui Chen, Vishal M Patel, Carlos D Castillo, and Rama Chellappa. Unconstrained still/video-based face verification with deep convolutional neural networks. *International Journal of Computer Vision*, 126(2-4):272–291, 2018. [2](#)
8. Sheng Chen, Yang Liu, Xiang Gao, and Zhen Han. Mobile-facenets: Efficient cnns for accurate real-time face verification on mobile devices. In *Chinese Conference on Biometric Recognition*, pages 428–438. Springer, 2018. [2](#)
9. Sumit Chopra, Raia Hadsell, Yann LeCun, et al. Learning a similarity metric discriminatively, with application to face verification. In *CVPR (1)*, pages 539–546, 2005. [1](#)

10. Aruni Roy Chowdhury, Tsung-Yu Lin, Subhransu Maji, and Erik Learned-Miller. One-to-many face recognition with bi-linear cnns. In *2016 IEEE Winter Conference on Applications of Computer Vision (WACV)*, pages 1–9. IEEE, 2016. 2
11. Nate Crosswhite, Jeffrey Byrne, Chris Stauffer, Omkar Parkhi, Qiong Cao, and Andrew Zisserman. Template adaptation for face verification and identification. *Image and Vision Computing*, 79:35–48, 2018. 1
12. Zhen Cui, Wen Li, Dong Xu, Shiguang Shan, and Xilin Chen. Fusing robust face region descriptors via multiple metric learning for face recognition in the wild. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3554–3561, 2013. 1
13. Jiankang Deng, Jia Guo, Niannan Xue, and Stefanos Zafeiriou. Arcface: Additive angular margin loss for deep face recognition. In *Proceedings of the IEEE Conference*
14. Changxing Ding and Dacheng Tao. Trunk-branch ensemble convolutional neural networks for video-based face recognition. *IEEE transactions on pattern analysis and machine intelligence*, 40(4):1002–1014, 2017. 2
15. Yin Fan, Xiangju Lu, Dian Li, and Yuanliu Liu. Video-based emotion recognition using cnn-rnn and c3d hybrid networks. In *Proceedings of the 18th ACM International Conference on Multimodal Interaction*, pages 445–450. ACM, 2016. 2
16. Sixue Gong, Yichun Shi, and Anil Jain. Low quality video face recognition: Multi-mode aggregation recurrent network (marn). In *Proceedings of the IEEE International Conference on Computer Vision Workshops*, pages 0–0, 2019. 2
17. Sixue Gong, Yichun Shi, and Anil K Jain. Video face recognition: Component-wise feature aggregation network (c-fan). *arXiv preprint arXiv:1902.07327*, 2019. 2, 3
18. Alex Graves, Christoph Mayer, Matthias Wimmer, Jürgen Schmidhuber, and Bernd Radig. Facial expression recognition with recurrent neural networks. In *Proceedings of the International Workshop on Cognition for Technical Systems*, 2008. 2
19. Jinwei Gu, Xiaodong Yang, Shalini De Mello, and Jan Kautz. Dynamic facial analysis: From bayesian filtering to recurrent neural network. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1548–1557, 2017. 2
20. Yandong Guo, Lei Zhang, Yuxiao Hu, Xiaodong He, and Jianfeng Gao. Ms-celeb-1m: A dataset and benchmark for large-scale face recognition. In *European Conference on Computer Vision*, pages 87–102. Springer, 2016. 2
21. Raia Hadsell, Sumit Chopra, and Yann 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. IEEE, 2006. 1
22. Mehrtash T Harandi, Conrad Sanderson, Sareh Shirazi, and Brian C Lovell. Graph embedding discriminant analysis on grassmannian manifolds for improved image set matching. In *CVPR 2011*, pages 2705–2712. IEEE, 2011. 2
23. Elad Hoffer and Nir Ailon. Deep metric learning using triplet network. In *International Workshop on Similarity-Based Pattern Recognition*, pages 84–92. Springer, 2015. 1
24. Junlin Hu, Jiwen Lu, and Yap-Peng Tan. Discriminative deep metric learning for face verification in the wild. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1875–1882, 2014. 1
25. Yiqun Hu, Ajmal S Mian, and Robyn Owens. Sparse approximated nearest points for image set classification. In *CVPR 2011*, pages 121–128. IEEE, 2011. 2
26. Zhiwu Huang, Ruiping Wang, Shiguang Shan, Xianqiu Li, and Xilin Chen. Log-euclidean metric learning on symmetric positive definite manifold with application to image set classification. In *International conference on machine learning*, pages 720–729, 2015. 2
27. Kuang-Chih Lee, Jeffrey Ho, Ming-Hsuan Yang, and David Kriegman. Video-based face recognition using probabilistic appearance manifolds. In *IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, volume 1, pages 1–313. Citeseer, 2003. 2
28. Haoxiang Li, Gang Hua, Zhe Lin, Jonathan Brandt, and Jianchao Yang. Probabilistic elastic matching for pose variant face verification. In *Proceedings of the IEEE Conference on Computer Vision and Pat-*

- tern Recognition*, pages 3499–3506, 2013. [1](#)
29. Liyuan Liu, Haoming Jiang, Pengcheng He, Weizhu Chen, Xiaodong Liu, Jianfeng Gao, and Jiawei Han. On the variance of the adaptive learning rate and beyond. *arXiv preprint arXiv:1908.03265*, 2019. [4](#)
 30. Luoqi Liu, Li Zhang, Hairong Liu, and Shuicheng Yan. Toward large-population face identification in unconstrained videos. *IEEE Transactions on Circuits and Systems for Video Technology*, 24(11):1874–1884, 2014. [1](#)
 31. Weiyang Liu, Yandong Wen, Zhiding Yu, Ming Li, Bhiksha Raj, and Le Song. SpheroFace: Deep hypersphere embedding for face recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 212–220, 2017. [1](#)
 32. Weiyang Liu, Yandong Wen, Zhiding Yu, and Meng Yang. Large-margin softmax loss for convolutional neural networks. In *Proceedings of The 33rd International Conference on Machine Learning*, pages 507–516, 2016. [1](#)
 33. Xiaofeng Liu, BVK Vijaya Kumar, Chao Yang, Qingming Tang, and Jane You. Dependency-aware attention control for unconstrained face recognition with image sets. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 548–565, 2018. [1](#)
 34. Jiwen Lu, Gang Wang, and Pierre Moulin. Image set classification using holistic multiple order statistics features and localized multi-kernel metric learning. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 329–336, 2013. [2](#)
 35. Brianna Maze, Jocelyn Adams, James A Duncan, Nathan Kalka, Tim Miller, Charles Otto, Anil K Jain, W Tyler Niggel, Janet Anderson, Jordan Cheney, et al. Iarpa janus benchmark-c: Face dataset and protocol. In *2018 International Conference on Biometrics (ICB)*, pages 158–165. IEEE, 2018. [3](#)
 36. Heydi Méndez-Vázquez, Yoanna Martínez-Díaz, and Zhenhua Chai. Volume structured ordinal features with background similarity measure for video face recognition. In *2013 International Conference on Biometrics (ICB)*, pages 1–6. IEEE, 2013. [1](#)
 37. Mostafa Parchami, Saman Bashbaghi, Eric Granger, and Saif Sayed. Using deep autoencoders to learn robust domain-invariant representations for still-to-video face recognition. In *2017 14th IEEE International Conference on Advanced Video and Signal Based Surveillance (AVSS)*, pages 1–6. IEEE, 2017. [2](#)
 38. Omkar M Parkhi, Karen Simonyan, Andrea Vedaldi, and Andrew Zisserman. A compact and discriminative face track descriptor. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1693–1700, 2014. [1](#)
 39. Rajeev Ranjan, Ankan Bansal, Hongyu Xu, Swami Sankaranarayanan, Jun-Cheng Chen, Carlos D Castillo, and Rama Chellappa. Crystal loss and quality pooling for unconstrained face verification and recognition. *arXiv preprint arXiv:1804.01159*, 2018. [2](#), [3](#)
 40. Yongming Rao, Jiwen Lu, and Jie Zhou. Attention-aware deep reinforcement learning for video face recognition. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 3931–3940, 2017. [1](#)
 41. Yongming Rao, Jiwen Lu, and Jie Zhou. Learning discriminative aggregation network for video-based face recognition and person re-identification. *International Journal of Computer Vision*, 127(6-7):701–718, 2019. [1](#)
 42. Florian Schroff, Dmitry Kalenichenko, and James Philbin. Facenet: A unified embedding for face recognition and clustering. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 815–823, 2015. [1](#)
 43. Gregory Shakhnarovich, John W Fisher, and Trevor Darrell. Face recognition from long-term observations. In *European Conference on Computer Vision*, pages 851–865. Springer, 2002. [2](#)
 44. Kihyuk Sohn, Sifei Liu, Guangyu Zhong, Xiang Yu, Ming-Hsuan Yang, and Manmohan Chandraker. Unsupervised domain adaptation for face recognition in unlabeled videos. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 3210–3218, 2017. [1](#)
 45. Yi Sun, Xiaogang Wang, and Xiaoou Tang. Deeply learned face representations are sparse, selective, and robust. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2892–2900, 2015. [1](#)
 46. Yaniv Taigman, Ming Yang, Marc’Aurelio Ranzato, and Lior Wolf. Deepface: Closing the gap to human-

- level performance in face verification. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1701–1708, 2014. [1](#)
47. Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *Advances in neural information processing systems*, pages 5998–6008, 2017. [2](#)
 48. Dayong Wang, Charles Otto, and Anil K Jain. Face search at scale. *IEEE transactions on pattern analysis and machine intelligence*, 39(6):1122–1136, 2016. [1](#)
 49. Feng Wang, Jian Cheng, Weiyang Liu, and Haijun Liu. Additive margin softmax for face verification. *IEEE Signal Processing Letters*, 25(7):926–930, 2018. [4](#)
 50. Hao Wang, Yitong Wang, Zheng Zhou, Xing Ji, Dihong Gong, Jingchao Zhou, Zhifeng Li, and Wei Liu. Cosface: Large margin cosine loss for deep face recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 5265–5274, 2018. [1](#)
 51. Jiang Wang, Yang Song, Thomas Leung, Chuck Rosenberg, Jingbin Wang, James Philbin, Bo Chen, and Ying Wu. Learning fine-grained image similarity with deep ranking. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1386–1393, 2014. [1](#)
 52. Ruiping Wang, Shiguang Shan, Xilin Chen, and Wen Gao. Manifold-manifold distance with application to face recognition based on image set. In *2008 IEEE Conference on Computer Vision and Pattern Recognition*, pages 1–8. IEEE, 2008. [2](#)
 53. Yandong Wen, Kaipeng Zhang, Zhifeng Li, and Yu Qiao. A discriminative feature learning approach for deep face recognition. In *European conference on computer vision*, pages 499–515. Springer, 2016. [1](#)
 54. Weidi Xie, Li Shen, and Andrew Zisserman. Comparator networks. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 782–797, 2018. [1](#)
 55. Jiaolong Yang, Peiran Ren, Dongqing Zhang, Dong Chen, Fang Wen, Hongdong Li, and Gang Hua. Neural aggregation network for video face recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4362–4371, 2017. [1](#), [2](#), [3](#)
 56. Jiaolong Yang, Peiran Ren, Dongqing Zhang, Dong Chen, Fang Wen, Hongdong Li, and Gang Hua. Neural aggregation network for video face recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4362–4371, 2017. [2](#)
 57. Meng Yang, Pengfei Zhu, Luc Van Gool, and Lei Zhang. Face recognition based on regularized nearest points between image sets. In *2013 10th IEEE international conference and workshops on automatic face and gesture recognition (FG)*, pages 1–7. IEEE, 2013. [2](#)
 58. Kaipeng Zhang, Zhanpeng Zhang, Zhifeng Li, and Yu Qiao. Joint face detection and alignment using multitask cascaded convolutional networks. *IEEE Signal Processing Letters*, 23(10):1499–1503, 2016. [1](#), [4](#)
 59. Tong Zhang, Wenming Zheng, Zhen Cui, Yuan Zong, and Yang Li. Spatial-temporal recurrent neural network for emotion recognition. *IEEE transactions on cybernetics*, 49(3):839–847, 2018. [2](#)