

Learning Unified Representations for Multi-Resolution Face Recognition

Hulingxiao He¹
hlxhe@bit.edu.cn

Wu Yuan¹
yuanwu@bit.edu.cn

Yidian Huang¹
hyd15213136303@gmail.com

Shilong Zhao¹
zhaoshilong0108@126.com

Wen Yuan*²
yuanw@lreis.ac.cn

Hanqing Li²
const.lhg@gmail.com

¹ Beijing Institute of Technology
Beijing, China

² Institute of Geographic Sciences and
Natural Resources Research, CAS
Beijing, China

Abstract

In this work, we propose Branch-to-Trunk network (BTNet), a representation learning method for multi-resolution face recognition. It consists of a trunk network (TNet), namely a unified encoder, and multiple branch networks (BNet), namely resolution adapters. As per the input, a resolution-specific BNet is used and the output are implanted as feature maps in the feature pyramid of TNet, at a layer with the same resolution. The discriminability of tiny faces is significantly improved, as the interpolation error introduced by rescaling, especially up-sampling, is mitigated on the inputs. With branch distillation and backward-compatible training, BTNet transfers discriminative high-resolution information to multiple branches while guaranteeing representation compatibility. Our experiments demonstrate strong performance on face recognition benchmarks, both for multi-resolution identity matching and feature aggregation, with much less computation amount and parameter storage. We establish new state-of-the-art on the challenging QMUL-SurvFace 1: N face identification task. Our code is available at <https://github.com/StevenSmith2000/BTNet>.

1 Introduction

Machine learning has made great strides with deep learning methods, but faces challenges with different types of data like structure and size. For example, face recognition

models struggle with changes in factors such as lighting and resolution when moving from the training domain to the testing domain.

Most face recognition methods map each image to a point embedding in the common metric space by deep neural networks (DNNs). The dissimilarity of images can be then calculated using various distance metrics (e.g., cosine similarity, Euclidean distance, etc.) for face recognition tasks. Face recognition models typically use deep neural networks to map images to a common metric space where the distance between two embeddings represents their dissimilarity. Recent advancements in margin-based loss [9] [67] [67] have improved the discriminability of the metric space, but lack of variation in training data can still lead to poor generalizability.

As known, the resolutions of face images in reality may be far beyond the scope covered by the model. As the small feature maps with a fixed spatial extent (e.g., 7×7) are mapped to an embedding with a predefined dimension (e.g., $128 - d$, $512 - d$, etc.) by a fully connected (fc) layer, input images need to be rescaled to a canonical spatial size (e.g., 112×112) before fed into the network. However, up-sampling low-resolution (LR) images introduces the interpolation error (see Section 3.1), deteriorating the recognizable ones which contain enough clues to identify the subject. Even though super-resolution methods ([2], [1], [68], [4], [68], [6], [68]) are widely used to build faces with good visualization, they inevitably introduce feature information of other identities when reconstructing high-resolution (HR) faces. This may lead to erroneous identity-specific features, which are detrimental to risk-controlled face recognition.

To improve discriminability while ensure the compatibility of the metric space for multi-resolution face representation, we learn the “unified” representation by a partially-coupled Branch-to-Trunk Network (BTNet). It is composed of multiple independent branch networks (BNet) and a shared trunk network (TNet). A resolution-specific BNet is used for a given image, and the output are implanted as feature maps in the feature pyramid of TNet, at a layer with the same resolution.

Furthermore, we find that multi-resolution training can be beneficial to building a strong and robust TNet, and backward-compatible training (BCT) [43] can improve the representation compatibility during the training process of BTNet. To ameliorate the discriminability of tiny faces, we propose branch distillation in intermediate layers, utilizing information extracted from HR images to help the extraction of discriminative features for resolution-specific branches.

Our method is simple and efficient, which can serve as a general framework easily applied to existing networks to improve their robustness against image resolutions. Since multi-resolution face recognition is dominated by super-resolution and projection methods, to the best of our knowledge, our method is the first attempt to decouple the information flow conditioned on the input resolution, which breaks the convention of up-sampling the inputs. Meanwhile, BTNet is able to reduce the number of FLOPS by operating the inputs without excessive up-sampling, and per-resolution storage cost by only storing the learned branches and resolution-aware BNs [67], while re-using the copy of the trunk model.

We demonstrate that our method performs comparably in various open-set face recognition tasks (1:1 face verification and 1: N face identification), while meaningfully reduces the redundant computation cost and parameter storage. In the challenging QMUL-SurvFace 1: N face identification task [9], we establish new state-of-the-art by outperforming state-of-the-art models. Furthermore, by avoiding the ill-posed problem (i.e., image up-sampling), our approach also effectively reduces the additional noise and uncertainty of the representation, which plays a key role in reliable risk-controlled face recognition.

2 Related Work

Compatible Representation Learning: The task of compatible representation learning aims at encoding features that are interoperable with the features extracted from other models. Shen et. al. [23] first formulated the problem of backward-compatible learning (BCT) and proposed to utilize the old classifier for compatible feature learning. Since the multi-model fashion benefits representation learning with lower computation, our idea of cross-resolution representation learning can be modeled similar to cross-model compatibility [0, 9, 32, 43, 54], as metric space alignment for different resolutions.

Knowledge Distillation and Transfer: The concept of knowledge distillation (KD) was first proposed by Hinton et. al. in [15], which can be summarized as employing a large parameter model (teacher) to supervise the learning of a small parameter model (student). Distillation from intermediate features [13, 14, 18, 20, 55, 37, 59, 60, 53, 61, 53] is widely adopted to enhance the effectiveness of knowledge transfer. However, due to the “dark knowledge” hidden in the intermediate layers, additional subtle design is often required to match and rescale intermediate features.

Low Resolution Face Recognition: Its task includes low resolution-to-low resolution (LR-to-LR) matching and low resolution-to-high resolution (LR-to-HR) matching [6]. The work can be divided into two categories [29]: (1) Super-resolution (SR) based methods aim to upscale LR images to construct HR images and use them for feature extraction [0, 11, 58, 45, 58, 52, 58]. (2) Projection-based methods aim to extract adequate representations in different domains and project them into a common feature space [28, 54, 54]. SR approaches are able to build faces with good visualization, but inevitably introduce feature information of other identities when reconstructing corresponding HR faces, thus introducing noise for identity-specific features.

3 Learning Specific-Shared Feature Transfer

Instead of rescaling the inputs to a canonical size, we build multiple resolution-specific branches (BNets) that are used to map inputs to intermediate features with the same resolution and a resolution-shared trunk (TNet) to map feature maps with different resolutions to a high-dimension embedding. We gain several important properties by doing so: (1) Processing inputs on its original resolution can diminish the inevitably introduced error via up-sampling or information loss via down-sampling, thus preserving the discriminability of visual information with different resolutions. (2) Information streams of different resolutions are encoded uniformly, thus enabling the representation compatibility, which is particularly beneficial to open-set face recognition considering that a compatible metric space is the prerequisite for computing similarity. (3) This also effectively reduce the computation for LR images by supplying computational resources conditioned on the input resolution.

3.1 Up-Sampling Error Analysis

Figure 1 illustrates the experimental estimation of interpolation error, whose upper bound increases with the decline of the image resolution. Note that the error soars up when the resolution drops below 32 approximately which can be viewed as LR face images, consistent with the tiny-object criterion [51].

The results show that: (1) inputs with a resolution higher than around 32 can be considered in the same HR domain, since the error information introduced by up-sampling via interpolation can be ignored to a certain extent; (2) inputs with a resolution lower than around




Figure 1: Estimated Error Upperbound. (bilinear interpolation, average value for over 100 images) with the change of image resolution relative to resolution 112.

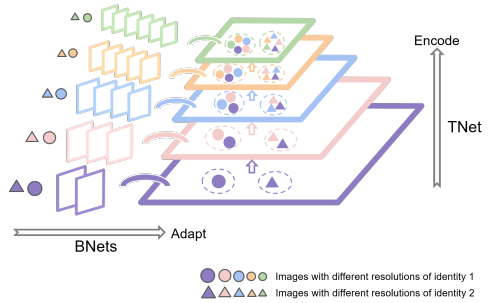


Figure 2: Basic ideas of the proposed BTNet. In this figure, feature maps with the same resolution are indicated by outlines in the same color.

32 should be treated as in various LR domains due to the high sensitivity of the resolution to errors.

3.2 Branch-to-Trunk Network

Let $X_{r'}$ be an input RGB image with a space shape: $X_{r'} \in \mathbb{R}^{H \times W \times 3}$, where $H \times W$ corresponds to the spatial dimension of the input and r' denotes the image resolution represented as $\min(H, W)$, $\max(H, W)$ or $\text{average}(H, W)$ based on the processing strategy. For efficient batch training and inference, we predefine a canonical size $S \times S$ (e.g., 112×112 for typical face recognition models like ArcFace [4]).

Figure 2 and 3 illustrate the main ideas of BTNet and an instantiation of BTNet framework, respectively. Our proposed **Branch-to-Trunk Network** (BTNet) consists of a trunk network $T: \mathbb{R}^{H \times W \times 3} \rightarrow \mathbb{R}^{C_{emb}}$ capable of extracting discriminative information with different resolutions and multiple branches B to focus on resolution-specific feature transfer independently. The work flow can be summarized as the following four steps: **1) Branch Selection:** input image $X_{r'}$ with resolution r' is first assigned with a resolution-specific branch B_r via the branch selection process to obtain X_r with resolution r , significantly reducing the scale of up-sampling compared to existing methods. **2) Resolution Adaptation:** the image $X_{r'}$ is encoded by the branch to obtain $z_r = B_r(X_{r'})$, which learns a mapping from the input image $X_{r'}$ to feature maps with the same resolution and expanded channels $z_r: \mathbb{R}^{r \times r \times 3} \rightarrow \mathbb{R}^{r \times r \times C_r}$. Note




Figure 3: Detailed architecture of BTNet-res50 (ϕ_{BT}). Note that ‘S’ and ‘U’ represent stage and unit respectively, and ‘/2’ means down-sampling by convolution with stride 2.

that C_r is predefined by the model design and doesn't depend on the resolution r' of the input image. Specifically, our branches B are implemented with same-resolution mapping: i.e., the model preserves the network architecture of T from input to the layer with resolution r and abandons down-sampling operations (e.g., replacing the convolution of stride 2 with stride 1, abandoning the pooling layers, etc.) to keep the same-resolution flow. **3) Unified Encoding:** The feature maps z_r are served as the input to the sub-network $T_r: \mathbb{R}^{r \times r \times C_r} \rightarrow \mathbb{R}^{C_{emb}}$ to obtain the final embedding $z_{final} = T_r(z_r)$; **4) Classification:** After obtaining the final embedding z_{final} of the input image, it is processed by fully connected layers to project to the probabilistic distribution for different identities.

3.3 Training Objectives

The training of BTNet includes training the trunk network T such that it can produce discriminative and compatible representations for multi-resolution information, and fine-tuning the branch networks B to encourage them to learn resolution-specific feature transfer, so as to improve accuracy without compromising compatibility.

Influence Loss. It is a compatibility-aware classification loss which is implemented by feeding the embeddings of the new model to the classifier of the old model [43]. There are various available loss functions that have been proven to be effective, like Triplet Loss [44], Center Loss [69], CosFace [66], Circle loss [48] et al. Thus, we can refine any loss function as our influence loss:

Any classification-based loss (e.g., NormFace [65], SphereFace [25], CosFace [66], ArcFace [4], etc.) can be refined as our influence loss. Since the difficulties of samples vary due to image resolution, we compute CurricularFace [47] as our classification loss in the original architecture, in the form of:

$$L_{influence} = L_{cur}(\varphi_{bt}, \kappa^*) \quad (1)$$

where φ_{bt} is the backbone (both B_r and T_r), and κ^* is the classifier of the pretrained trunk T .

Branch Distillation Loss. Due to the continuity of the scale change of both the image pyramid and the feature pyramid [24], we can get a qualitative sense of the similarity between images and feature maps with the same resolution (see Figure 4). Furthermore, features extracted from HR images have richer and clearer information than those from LR images [40]. Motivated by these analyses, we utilize an MSE loss to encourage the branch output z_r to be similar to the corresponding feature maps of the pretrained trunk network z_s :

$$L_{branch} = \frac{1}{V} \sum_{v=1}^V (z_{r_v} - z_{s_v})^2 \quad (2)$$

where V denotes the batch size.

The whole training objective is a combination of the above objectives:

$$L = L_{influence} + \lambda_{branch} L_{branch} \quad (3)$$

where λ_{branch} is a hyper-parameter to weigh the losses and we set $\lambda_{branch} = 0.5$ in all our experiments.




Figure 4: Visual comparison of face image-feature map pairs with different resolutions (resized to a common size here for illustration).




Figure 5: Comparison of # Params (M) between fully finetuning and φ_{br} .

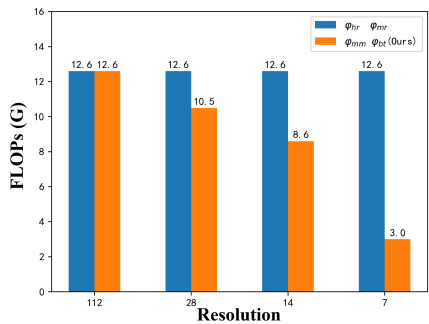


Figure 6: Comparison of FLOPs (G) between baselines and φ_{br} .

3.4 Storing Branch Networks

An obvious adaptation strategy is fully finetuning of the model on each resolution. However, this strategy requires one to store and deploy a separate copy of the backbone parameters for every resolution, which is an expensive proposition and difficult to expand into more segmented resolution branches. Our BTNet is beneficial in the scenario of multi-resolution face recognition which achieves better parameter/accuracy trade-offs. Since activation statistics including means and variances under different resolutions are incompatible [52], we update and store Batch Normalization (BN) [49] parameters in all layers of B_r and T_r for each resolution, whose amount is negligible. Apart from this, we only need to store the learned branches and re-use the original copy of the pretrained trunk model, significantly reducing the storage cost. Figure 5 shows that BTNet requires only 1.1% ~ 48.9% of all the parameters compared to fully updating all the parameters of TNet.

4 Experiments

To validate BTNet on face recognition tasks in open universe, we perform 1:1 verification and 1 : N identification tasks in two different settings, including (a) multi-resolution identity matching, and (b) multi-resolution feature aggregation.

4.1 Implementation Details

Datasets. We use MS1Mv3 [5] for training face embedding models. The MS1Mv3 dataset contains 5,179,510 images of 93,431 celebrities. We try on six widely adopted face verification benchmarks: LFW [16], CFP-FF [4], CFP-FP [24], AgeDB-30 [53], CALFW [66], and CPLFW [65], while the large-scale surveillance face dataset QMUL-SurvFace [8] is used for 1:N face identification, which contains native LR surveillance faces across wide space and time. The spatial resolution for QMUL-SurvFace ranges from 6/5 to 124/106 in height/width with an average of 24/20.

Baselines. In our experiment, several baselines are used to validate BTNet in learning discriminative and compatible representations for multi-resolution face recognition.

- High-Resolution Trained φ_{hr} .** Naive baseline trained with HR data.

- Independently Trained φ_{mm} .** Multi-model fashion: is it possible to achieve better results if we train a specific model for each resolution independently? Specifically, we train φ_r for data with resolution r and denote the multi-model collections as φ_{mm} .

Table 1: Comparison of different methods on six face verification benchmarks.

	Cross-resolution identity matching						Same-resolution identity matching							
	112&7		112&14		112&28		7&7		14&14		28&28		112&112	
	Acc.	Gain	Acc.	Gain	Acc.	Gain	Acc.	Gain	Acc.	Gain	Acc.	Gain	Acc.	Gain
ϕ_{hr}	57.75	-	81.02	-	95.90	-	60.70	-	73.88	-	93.58	-	97.68	-
ϕ_{mm}	50.58	-0.89	49.90	-4.82	50.03	-305.80	62.57	+1.00	78.00	+1.00	94.68	+1.00	97.68	-
ϕ_{mr}	65.85	+1.00	87.47	+1.00	96.05	+1.00	61.02	+0.17	80.32	+1.56	95.12	+1.40	97.25	-
$\phi_{mr(v2)}$	65.68	+0.98	87.13	+0.95	95.70	-1.33	60.82	+0.06	80.22	+1.54	95.63	+1.86	96.82	-
$\phi_{mr(v3)}$	68.80	+1.36	88.13	+1.10	96.62	+4.80	61.62	+0.49	80.55	+1.62	94.78	+1.09	97.52	-
ϕ_{hr} (Ours)	86.10	+3.50	94.08	+2.02	96.65	+5.00	77.78	+9.13	90.90	+4.13	96.27	+2.45	97.25	-

•**Multi-Resolution Trained ϕ_{mr} .** Trained with multi-resolution data which adapts to resolution-variance. For a comprehensive evaluation, we implemented three baselines, denoted as ϕ_{mr} , $\phi_{mr(v2)}$, $\phi_{mr(v3)}$ respectively. Each image is down-sampled to a certain size and then up-sampled to 112×112 . The differences are as follows: (i) ϕ_{mr} : down-sampled to a size in the candidate set $\{\frac{112}{2^i} \times \frac{112}{2^i} | i = 0, 1, 2, 3, 4\}$ with equal probability of being chosen. (ii) $\phi_{mr(v2)}$: down-sampled to a size in the candidate set with unequal probability [0.3 0.25 0.2 0.15 0.1]. (iii) $\phi_{mr(v3)}$: down-sampled to a size in the candidate interval [4, 112].

Instantiation of Network Architecture. The BTNet and baselines are implemented with ResNet50 [12], and they could be extended easily with other implementations.

4.2 Evaluation Metrics

On the benchmarks for face verification, we use 1:1 verification accuracy as the basic metrics. The rank-20 true positive identification rates (TPIR20) at varying false positive identification rates (FPIR) and AUC are used to report the identification results on QMUL-SurvFace.

For better evaluation, we define another two metrics to assess the relative performance gain similar to [52, 43].

Cross-Resolution Gain. With the purpose towards the cross-resolution compatible representations, we define the performance gain as follows:

$$Gain_{r_1 \& r_2}(\varphi) = \frac{M_{r_1 \& r_2}(\varphi) - M_{r_1 \& r_2}(\varphi_{hr})}{|M_{r_1 \& r_2}(\varphi_{mr}) - M_{r_1 \& r_2}(\varphi_{hr})|} \quad (4)$$

Here $M_{r_1 \& r_2}(\cdot)$ are metrics when the resolutions of the image/template pair are $r_1 \times r_1$ and $r_2 \times r_2$ ($r_1 \neq r_2$), respectively. φ_{mr} shares the same architecture with φ_{hr} while is trained on multi-resolution images and thus serves as the baseline of cross-resolution gain.

Same-Resolution Gain. For the scenario of multi-resolution face recognition, the performance of same-resolution verification/identification is also vital besides cross-resolution one. Therefore, we report the relative performance improvement from base model φ_{hr} in the scenario of same-resolution.

$$Gain_{r \& r}(\varphi) = \frac{M_{r \& r}(\varphi) - M_{r \& r}(\varphi_{hr})}{|M_{r \& r}(\varphi_r) - M_{r \& r}(\varphi_{hr})|} \quad (5)$$

Here $M_{r \& r}(\cdot)$ are metrics when the resolutions of the image/template pair are both $r \times r$. φ_r is a model of the set $\{\varphi_{mm} = \varphi | r = 7, 14, 28\}$ trained on images with resolution $r \times r$ without considering cross-resolution representation compatibility, which serves as the baseline of same-resolution gain on resolution r .

4.3 Results

4.3.1 Multi-Resolution Face Verification

We now conduct experiments on the proposed BTNet framework for multi-resolution identity matching. Two different settings are included : (1) same-resolution matching, and (2) cross-resolution matching. Table 1 compares the average performance on popular benchmarks for φ_{hr} , φ_{mm} , φ_{mr} , φ_{bt} .

When directly applied to test data with the resolution lower than training data, φ_{hr} suffers a severe performance degradation. Up-sampling images via interpolation can increase the amount of data but not the amount of information, only to improve the detailed part of the image and the spatial resolution (size) [26]. Moreover, it also brings various noise and artificial processing traces [46]. Up-sampling images via interpolation—typically bilinear interpolation or bicubic interpolation of 4x4 pixel neighborhoods, essentially a function approximation method, is bound to introduce error information, thus potentially confusing identity information, which is especially crucial for LR images with limited details. We are able to observe improvement of φ_{mm} in same-resolution matching but its cross-resolution gain is negative with approximately 50% accuracy. Unsurprisingly, independently trained φ_r is unaware of representation compatibility, and thus does not naturally suit for cross-resolution recognition. The results show that φ_{mr} improved both cross-resolution and same-resolution accuracy by a large margin, as it learns to adapt to resolution variance and maintain discriminability of multi-resolution inputs. Note that the model size and training data scale stay the same, while only the resolution distribution of the data changes for φ_{mr} , and thus there is a marginal accuracy drop in the setting of 112&112 matching. Comparably, φ_{bt} substantially outperforms all baselines with 2.02~5.00 cross-resolution gain and 2.45~9.13 same-resolution gain. Importantly, due to the multi-resolution branches, our approach has a cost same with φ_{mm} , significantly lower than φ_{hr} and φ_{mr} (see Figure 6).

4.3.2 Multi-Resolution Face Identification

In the native scenario, it is common to inference on inputs with resolutions not strictly matched to the branch. Since the low-quality image may possess an underlying optical resolution significantly lower than its size due to degraded quality caused by noise, blur, occlusion, etc [60], there exists dislocation between the underlying optical resolution of native face images and that of a branch. To avoid introducing extra large-scale parameters for predicting the image quality, three heuristic selection strategies based on different resolution indicators are validated. Table 2 compares BTNet against the state-of-the-arts models

Table 2: Performance of face identification on QMUL-SurvFace.

	TPiR20(%)@FPIR			
	AUC	0.3	0.2	0.1
VGG-Face [45]	14.0	5.1	2.6	0.8
DeepID2 [45]	20.8	12.8	8.1	3.4
FaceNet [45]	19.8	12.7	8.1	4.3
SphereFace [45]	28.1	21.3	15.7	8.3
SRCNN [45]	27.0	20.0	14.9	6.2
FSRCNN [45]	27.3	20.0	14.4	6.1
VDSR [45]	27.3	20.1	14.5	6.1
DRRN [45]	27.5	20.3	14.9	6.3
LapSRN [45]	27.4	20.2	14.7	6.3
ArcFace [45]	25.3	18.7	15.1	10.1
RAN [45]	32.3	26.5	21.6	14.9
SST [45]	-	12.4	-	9.7
MASST [45]	-	12.2	-	9.2
MIND-Net [45]	31.9	25.5	-	20.4
AdaFace [45]	32.6	28.3	23.6	16.5
BTNet (avg.+floor)	32.6	27.9	23.4	16.5
BTNet (avg.+near)	34.6	30.3	25.7	18.9
BTNet (avg.+ceil)	35.4	31.1	26.8	20.3
BTNet (min.+floor)	32.3	27.6	23.2	16.1
BTNet (min.+near)	34.0	29.6	25.0	18.0
BTNet (min.+ceil)	35.3	31.0	26.6	19.9
BTNet (max.+floor)	33.6	29.1	24.5	17.6
BTNet (max.+near)	35.2	31.0	26.4	19.6
BTNet (max.+ceil)	35.4	31.2	26.9	20.6

Table 3: Comparison of different training methods for our BTNet.

Training method	Acc. (%)		# Params.
	112&14	14&14	
Scratch	49.90	78.00	43.59M
Pretraining	78.05	76.87	43.59M
Pretraining + BCT	85.90	78.04	43.59M
Pretraining + BCT + Fix Trunk	85.07	77.22	2.29M
Pretraining + BCT + Fix Trunk + Branch Distillation	94.08	90.90	2.29M

Table 4: Ablation study of different loss functions.

Implementation of influence loss	112&14 Acc.(%)	14&14 Acc.(%)
CosFace	94.10	90.78
ArcFace	94.17	90.88
CurricularFace	94.08	90.90

on QMUL-SurvFace 1:N identification benchmark. We are able to observe that our proposed approach extends the state-of-the-arts while being more computationally efficient. We believe the performance of BTNet (max + ceil) is the highest that have been reported so far, and we believe it is meaningful with the increased focus on unconstrained surveillance applications.

5 Ablation Study

In all these experiments, we report the average verification results on six benchmarks in 112&14 and 14&14 matching, representing cross-resolution and same-resolution performance respectively.

Training Method Alternatives. Here, we experimentally compare different training methods: (1) Scratch: train without pretrained trunk parameters. (2) Pretraining: initialize the backbone and classifier with the pretrained trunk network. (3) Backward-compatible training (BCT [43]): fix parameters of the old classifier. (4) Fix-trunk: fix parameters of the trunk subnet T_r . (5) Branch distillation: use L2-distance to obtain the loss between the intermediate feature maps at the coupling layer of the pretrained trunk T and the branch B_r .

We compare different training method combinations in Table 3 and find that both pretraining and BCT succeeded in ensuring representation compatibility. Among these two, BCT performs better since it imposes a stricter constraint during training. Furthermore, we are able to observe that branch distillation is crucial for improving the discriminative power by transferring high-resolution information to low-resolution branches.

Loss Functions. Since the difficulties of samples vary due to image resolution, we compute CurricularFace [44] as our classification loss in the original architecture, which distinguishes both the difficultness of different samples in each stage and relative importance of easy and hard samples during different training stages.

To prove the main technical contribution of BTNet (rather than other components), we use different loss functions to replace the CurricularFace loss as influence loss in the original architecture. The comparison results(in Table 4) demonstrate that there is no significant difference among different implementations of influence loss. It means that the main performance gain is attributed to our design.

Where should we have resolution-specific layers? We conducted an ablation to see the effects of different specific-shared layer allocation strategies. The experiment was done with different trunk layers (i.e., the parameters of these layers are inherited from the pretrained trunk without updating). Figure 7 shows the results. We find that increasing the number of branch layers (i.e., specific layers for different resolutions) will lead to better performance due to increased flexibility. Our specific-shared layer allocation of BTNet can achieve better

parameter/accuracy tradeoffs as further increasing the number of trunk layers based on BT-Net cannot lead to significantly better performance but increases parameter storage cost by a large margin.

6 Discussion and Conclusion

This paper works on the problem of multi-resolution face recognition, and provides a new scheme to operate images conditioned on its input resolution without large span rescaling. The error introduced by up-sampling via interpolation is investigated and analyzed. Decoupled as branches for discriminative representation learning and coupled as the trunk for compatible representation learning, our Branch-to-Trunk Network (BTNet) achieves significant improvements on multi-resolution face verification and identification tasks. Besides, the superiority of BTNet in reducing computational cost and parameter storage cost is also demonstrated.

Limitations and Future Work. The dislocation between the underlying optical resolution of native face images and that of a certain branch may limit the power of the model, which may be improved by selecting the optimal processing branch for the input in combination with the image quality, rather than by image size alone. In the experiments, we provide an intuitive way to select the branch for inputs (see Figure 8). Importantly, based on the unified multi-resolution metric space, the underlying resolution of the inputs (integrated spatial resolution with quality assessment) can be utilized to provide the reliability of the representation and contribute to risk-controlled face recognition. They will be our future research directions.




Figure 8: Branch selection process. Max/min/average is used on (W, H) to obtain a resolution indicator for further allocation (floor/near/ceil) to a certain branch.




Figure 7: Comparison of verification accuracy and the amount of stored parameters for different specific-shared layer allocation strategies.

Acknowledgements. This work is partly supported by Strategic Priority Research Program of the Chinese Academy of Sciences, No.XDA23100103 and National Key R&D Program of China (No. 2022YFF0711601).

References

- [1] Mateusz Budnik and Yannis Avrithis. Asymmetric metric learning for knowledge transfer. In *IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2021, virtual, June 19-25, 2021*, pages 8228–8238. Computer Vision Foundation / IEEE, 2021. URL https://openaccess.thecvf.com/content/CVPR2021/html/Budnik_Asymmetric_Metric_Learning_for_Knowledge_Transfer_CVPR_2021_paper.html.
- [2] Zhiyi Cheng, Xiatian Zhu, and Shaogang Gong. Low-resolution face recognition. In C. V. Jawahar, Hongdong Li, Greg Mori, and Konrad Schindler, editors, *Computer Vision - ACCV 2018 - 14th Asian Conference on Computer Vision, Perth, Australia, December 2-6, 2018, Revised Selected Papers, Part III*, volume 11363 of *Lecture Notes in Computer Science*, pages 605–621. Springer, 2018. doi: 10.1007/978-3-030-20893-6_38. URL https://doi.org/10.1007/978-3-030-20893-6_38.
- [3] Zhiyi Cheng, Xiatian Zhu, and Shaogang Gong. Surveillance face recognition challenge. *CoRR*, abs/1804.09691, 2018. URL <http://arxiv.org/abs/1804.09691>.
- [4] Jiankang Deng, Jia Guo, Niannan Xue, and Stefanos Zafeiriou. Arcface: Additive angular margin loss for deep face recognition. In *IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2019, Long Beach, CA, USA, June 16-20, 2019*, pages 4690–4699. Computer Vision Foundation / IEEE, 2019. doi: 10.1109/CVPR.2019.00482. URL http://openaccess.thecvf.com/content_CVPR_2019/html/Deng_ArcFace_Additive_Angular_Margin_Loss_for_Deep_Face_Recognition_CVPR_2019_paper.html.
- [5] Jiankang Deng, Jia Guo, Debing Zhang, Yafeng Deng, Xiangju Lu, and Song Shi. Lightweight face recognition challenge. In *2019 IEEE/CVF International Conference on Computer Vision Workshops, ICCV Workshops 2019, Seoul, Korea (South), October 27-28, 2019*, pages 2638–2646. IEEE, 2019. doi: 10.1109/ICCVW.2019.00322. URL <https://doi.org/10.1109/ICCVW.2019.00322>.
- [6] Chao Dong, Chen Change Loy, Kaiming He, and Xiaoou Tang. Learning a deep convolutional network for image super-resolution. In David J. Fleet, Tomás Pa-jdla, Bernt Schiele, and Tinne Tuytelaars, editors, *Computer Vision - ECCV 2014 - 13th European Conference, Zurich, Switzerland, September 6-12, 2014, Proceedings, Part IV*, volume 8692 of *Lecture Notes in Computer Science*, pages 184–199. Springer, 2014. doi: 10.1007/978-3-319-10593-2_13. URL https://doi.org/10.1007/978-3-319-10593-2_13.
- [7] Chao Dong, Chen Change Loy, and Xiaoou Tang. Accelerating the super-resolution convolutional neural network. In Bastian Leibe, Jiri Matas, Nicu Sebe, and Max Welling, editors, *Computer Vision - ECCV 2016 - 14th European Conference, Amsterdam, The Netherlands, October 11-14, 2016, Proceedings, Part II*, volume 9906 of *Lecture Notes in Computer Science*, pages 391–407. Springer, 2016. doi: 10.1007/978-3-319-46475-6_25. URL https://doi.org/10.1007/978-3-319-46475-6_25.

- [8] Hang Du, Hailin Shi, Yuchi Liu, Jun Wang, Zhen Lei, Dan Zeng, and Tao Mei. Semi-siamese training for shallow face learning. In *European Conference on Computer Vision*, pages 36–53. Springer, 2020.
- [9] Rahul Duggal, Hao Zhou, Shuo Yang, Yuanjun Xiong, Wei Xia, Zhuowen Tu, and Stefano Soatto. Compatibility-aware heterogeneous visual search. In *IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2021, virtual, June 19-25, 2021*, pages 10723–10732. Computer Vision Foundation / IEEE, 2021. URL https://openaccess.thecvf.com/content/CVPR2021/html/Duggal_Compatibility-Aware_Heterogeneous_Visual_Search_CVPR_2021_paper.html.
- [10] Han Fang, Weihong Deng, Yaoyao Zhong, and Jiani Hu. Generate to adapt: Resolution adaptation network for surveillance face recognition. In Andrea Vedaldi, Horst Bischof, Thomas Brox, and Jan-Michael Frahm, editors, *Computer Vision - ECCV 2020 - 16th European Conference, Glasgow, UK, August 23-28, 2020, Proceedings, Part XV*, volume 12360 of *Lecture Notes in Computer Science*, pages 741–758. Springer, 2020. doi: 10.1007/978-3-030-58555-6_44. URL https://doi.org/10.1007/978-3-030-58555-6_44.
- [11] Klemen Grm, Walter J. Scheirer, and Vitomir Struc. Face hallucination using cascaded super-resolution and identity priors. *IEEE Trans. Image Process.*, 29:2150–2165, 2020. doi: 10.1109/TIP.2019.2945835. URL <https://doi.org/10.1109/TIP.2019.2945835>.
- [12] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *2016 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2016, Las Vegas, NV, USA, June 27-30, 2016*, pages 770–778. IEEE Computer Society, 2016. doi: 10.1109/CVPR.2016.90. URL <https://doi.org/10.1109/CVPR.2016.90>.
- [13] Byeongho Heo, Jeesoo Kim, Sangdoon Yun, Hyojin Park, Nojun Kwak, and Jin Young Choi. A comprehensive overhaul of feature distillation. In *2019 IEEE/CVF International Conference on Computer Vision, ICCV 2019, Seoul, Korea (South), October 27 - November 2, 2019*, pages 1921–1930. IEEE, 2019. doi: 10.1109/ICCV.2019.00201. URL <https://doi.org/10.1109/ICCV.2019.00201>.
- [14] Byeongho Heo, Minsik Lee, Sangdoon Yun, and Jin Young Choi. Knowledge transfer via distillation of activation boundaries formed by hidden neurons. In *The Thirty-Third AAAI Conference on Artificial Intelligence, AAAI 2019, The Thirty-First Innovative Applications of Artificial Intelligence Conference, IAAI 2019, The Ninth AAAI Symposium on Educational Advances in Artificial Intelligence, EAAI 2019, Honolulu, Hawaii, USA, January 27 - February 1, 2019*, pages 3779–3787. AAAI Press, 2019. doi: 10.1609/aaai.v33i01.33013779. URL <https://doi.org/10.1609/aaai.v33i01.33013779>.
- [15] Geoffrey E. Hinton, Oriol Vinyals, and Jeffrey Dean. Distilling the knowledge in a neural network. *CoRR*, abs/1503.02531, 2015. URL <http://arxiv.org/abs/1503.02531>.

- [16] Gary B Huang, Marwan Mattar, Tamara Berg, and Eric Learned-Miller. Labeled faces in the wild: A database for studying face recognition in unconstrained environments. In *Workshop on faces in 'Real-Life' Images: detection, alignment, and recognition*, 2008.
- [17] Yuge Huang, Yuhan Wang, Ying Tai, Xiaoming Liu, Pengcheng Shen, Shaoxin Li, Jilin Li, and Feiyue Huang. Curricularface: Adaptive curriculum learning loss for deep face recognition. In *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2020, Seattle, WA, USA, June 13-19, 2020*, pages 5900–5909. Computer Vision Foundation / IEEE, 2020. doi: 10.1109/CVPR42600.2020.00594. URL https://openaccess.thecvf.com/content_CVPR_2020/html/Huang_CurricularFace_Adaptive_Curriculum_Learning_Loss_for_Deep_Face_Recognition_CVPR_2020_paper.html.
- [18] Zehao Huang and Naiyan Wang. Like what you like: Knowledge distill via neuron selectivity transfer. *CoRR*, abs/1707.01219, 2017. URL <http://arxiv.org/abs/1707.01219>.
- [19] Sergey Ioffe and Christian Szegedy. Batch normalization: Accelerating deep network training by reducing internal covariate shift. In Francis R. Bach and David M. Blei, editors, *Proceedings of the 32nd International Conference on Machine Learning, ICML 2015, Lille, France, 6-11 July 2015*, volume 37 of *JMLR Workshop and Conference Proceedings*, pages 448–456. JMLR.org, 2015. URL <http://proceedings.mlr.press/v37/ioffe15.html>.
- [20] Jangho Kim, Seonguk Park, and Nojun Kwak. Paraphrasing complex network: Network compression via factor transfer. In Samy Bengio, Hanna M. Wallach, Hugo Larochelle, Kristen Grauman, Nicolò Cesa-Bianchi, and Roman Garnett, editors, *Advances in Neural Information Processing Systems 31: Annual Conference on Neural Information Processing Systems 2018, NeurIPS 2018, December 3-8, 2018, Montréal, Canada*, pages 2765–2774, 2018. URL <https://proceedings.neurips.cc/paper/2018/hash/6d9cb7de5e8ac30bd5e8734bc96a35c1-Abstract.html>.
- [21] Jiwon Kim, Jung Kwon Lee, and Kyoung Mu Lee. Accurate image super-resolution using very deep convolutional networks. In *2016 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2016, Las Vegas, NV, USA, June 27-30, 2016*, pages 1646–1654. IEEE Computer Society, 2016. doi: 10.1109/CVPR.2016.182. URL <https://doi.org/10.1109/CVPR.2016.182>.
- [22] Minchul Kim, Anil K Jain, and Xiaoming Liu. Adaface: Quality adaptive margin for face recognition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 18750–18759, 2022.
- [23] Wei-Sheng Lai, Jia-Bin Huang, Narendra Ahuja, and Ming-Hsuan Yang. Deep laplacian pyramid networks for fast and accurate super-resolution. In *2017 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2017, Honolulu, HI, USA, July 21-26, 2017*, pages 5835–5843. IEEE Computer Society, 2017. doi: 10.1109/CVPR.2017.618. URL <https://doi.org/10.1109/CVPR.2017.618>.
- [24] Tony Lindeberg. Scale-space theory: a basic tool for analyzing structures at different scales. *Journal of Applied Statistics*, 21(1-2):225–270, 1994. doi: 10.1080/757582976. URL <https://doi.org/10.1080/757582976>.

- [25] Weiyang Liu, Yandong Wen, Zhiding Yu, Ming Li, Bhiksha Raj, and Le Song. Spheraface: Deep hypersphere embedding for face recognition. In *2017 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2017, Honolulu, HI, USA, July 21-26, 2017*, pages 6738–6746. IEEE Computer Society, 2017. doi: 10.1109/CVPR.2017.713. URL <https://doi.org/10.1109/CVPR.2017.713>.
- [26] Z. G. Liu and D. Z. Liu. Reappraising about image magnification methods based on wavelet transformation. *Journal of Image and Graphics*, 2003.
- [27] Cheng-Yaw Low, Andrew Beng-Jin Teoh, and Jaewoo Park. Mind-net: A deep mutual information distillation network for realistic low-resolution face recognition. *IEEE Signal Processing Letters*, 28:354–358, 2021.
- [28] Ze Lu, Xudong Jiang, and Alex C. Kot. Deep coupled resnet for low-resolution face recognition. *IEEE Signal Process. Lett.*, 25(4):526–530, 2018. doi: 10.1109/LSP.2018.2810121. URL <https://doi.org/10.1109/LSP.2018.2810121>.
- [29] Luis S. Luevano, Leonardo Chang, Heydi Méndez-Vázquez, Yoanna Martínez-Díaz, and Miguel González-Mendoza. A study on the performance of unconstrained very low resolution face recognition: Analyzing current trends and new research directions. *IEEE Access*, 9:75470–75493, 2021. doi: 10.1109/ACCESS.2021.3080712. URL <https://doi.org/10.1109/ACCESS.2021.3080712>.
- [30] Yui Man Lui, David Bolme, Bruce A. Draper, J. Ross Beveridge, Geoff Givens, and P. Jonathon Phillips. A meta-analysis of face recognition covariates. In *2009 IEEE 3rd International Conference on Biometrics: Theory, Applications, and Systems*, pages 1–8, 2009. doi: 10.1109/BTAS.2009.5339025.
- [31] Yoanna Martínez-Díaz, Heydi Méndez-Vázquez, Luis S. Luevano, Leonardo Chang, and Miguel González-Mendoza. Lightweight low-resolution face recognition for surveillance applications. In *25th International Conference on Pattern Recognition, ICPR 2020, Virtual Event / Milan, Italy, January 10-15, 2021*, pages 5421–5428. IEEE, 2020. doi: 10.1109/ICPR48806.2021.9412280. URL <https://doi.org/10.1109/ICPR48806.2021.9412280>.
- [32] Qiang Meng, Chixiang Zhang, Xiaoqiang Xu, and Feng Zhou. Learning compatible embeddings. In *2021 IEEE/CVF International Conference on Computer Vision, ICCV 2021, Montreal, QC, Canada, October 10-17, 2021*, pages 9919–9928. IEEE, 2021. doi: 10.1109/ICCV48922.2021.00979. URL <https://doi.org/10.1109/ICCV48922.2021.00979>.
- [33] Stylianos Moschoglou, Athanasios Papaioannou, Christos Sagonas, Jiankang Deng, Irene Kotsia, and Stefanos Zafeiriou. Agedb: The first manually collected, in-the-wild age database. In *2017 IEEE Conference on Computer Vision and Pattern Recognition Workshops, CVPR Workshops 2017, Honolulu, HI, USA, July 21-26, 2017*, pages 1997–2005. IEEE Computer Society, 2017. doi: 10.1109/CVPRW.2017.250. URL <https://doi.org/10.1109/CVPRW.2017.250>.
- [34] Sivaram Prasad Mudunuri, Soubhik Sanyal, and Soma Biswas. Genlr-net: Deep framework for very low resolution face and object recognition with generalization to unseen categories. In *2018 IEEE Conference on Computer Vision and Pattern*

- Recognition Workshops, CVPR Workshops 2018, Salt Lake City, UT, USA, June 18-22, 2018*, pages 489–498. Computer Vision Foundation / IEEE Computer Society, 2018. doi: 10.1109/CVPRW.2018.00090. URL http://openaccess.thecvf.com/content_cvpr_2018_workshops/w11/html/Mudunuri_GenLR-Net_Deep_Framework_CVPR_2018_paper.html.
- [35] Wonpyo Park, Dongju Kim, Yan Lu, and Minsu Cho. Relational knowledge distillation. In *IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2019, Long Beach, CA, USA, June 16-20, 2019*, pages 3967–3976. Computer Vision Foundation / IEEE, 2019. doi: 10.1109/CVPR.2019.00409. URL http://openaccess.thecvf.com/content_cvpr_2019/html/Park_Relational_Knowledge_Distillation_CVPR_2019_paper.html.
- [36] Omkar M. Parkhi, Andrea Vedaldi, and Andrew Zisserman. Deep face recognition. In Xianghua Xie, Mark W. Jones, and Gary K. L. Tam, editors, *Proceedings of the British Machine Vision Conference 2015, BMVC 2015, Swansea, UK, September 7-10, 2015*, pages 41.1–41.12. BMVA Press, 2015. doi: 10.5244/C.29.41. URL <https://doi.org/10.5244/C.29.41>.
- [37] Baoyun Peng, Xiao Jin, Dongsheng Li, Shunfeng Zhou, Yichao Wu, Jiaheng Liu, Zhaoning Zhang, and Yu Liu. Correlation congruence for knowledge distillation. In *2019 IEEE/CVF International Conference on Computer Vision, ICCV 2019, Seoul, Korea (South), October 27 - November 2, 2019*, pages 5006–5015. IEEE, 2019. doi: 10.1109/ICCV.2019.00511. URL <https://doi.org/10.1109/ICCV.2019.00511>.
- [38] Aashish Rai, Vishal M. Chudasama, Kishor P. Upla, Kiran B. Raja, Raghavendra Ramachandra, and Christoph Busch. Comsupresnet: A compact super-resolution network for low-resolution face images. In *8th International Workshop on Biometrics and Forensics, IWBf 2020, Porto, Portugal, April 29-30, 2020*, pages 1–6. IEEE, 2020. doi: 10.1109/IWBf49977.2020.9107946. URL <https://doi.org/10.1109/IWBf49977.2020.9107946>.
- [39] Adriana Romero, Nicolas Ballas, Samira Ebrahimi Kahou, Antoine Chassang, Carlo Gatta, and Yoshua Bengio. Fitnets: Hints for thin deep nets. In Yoshua Bengio and Yann LeCun, editors, *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*, 2015. URL <http://arxiv.org/abs/1412.6550>.
- [40] Florian Schroff, Dmitry Kalenichenko, and James Philbin. Facenet: A unified embedding for face recognition and clustering. In *IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2015, Boston, MA, USA, June 7-12, 2015*, pages 815–823. IEEE Computer Society, 2015. doi: 10.1109/CVPR.2015.7298682. URL <https://doi.org/10.1109/CVPR.2015.7298682>.
- [41] 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.
- [42] Soumyadip Sengupta, Jun-Cheng Chen, Carlos Castillo, Vishal M Patel, Rama Chelappa, and David W Jacobs. Frontal to profile face verification in the wild. In *2016*

- IEEE winter conference on applications of computer vision (WACV)*, pages 1–9. IEEE, 2016.
- [43] Yantao Shen, Yuanjun Xiong, Wei Xia, and Stefano Soatto. Towards backward-compatible representation learning. In *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2020, Seattle, WA, USA, June 13-19, 2020*, pages 6367–6376. Computer Vision Foundation / IEEE, 2020. doi: 10.1109/CVPR42600.2020.00640. URL https://openaccess.thecvf.com/content_CVPR_2020/html/Shen_Towards_Backward-Compatible_Representation_Learning_CVPR_2020_paper.html.
- [44] Hailin Shi, Dan Zeng, Yichun Tai, Hang Du, Yibo Hu, and Tao Mei. Multi-agent semi-siamese training for long-tail and shallow face learning. *arXiv preprint arXiv:2105.04113*, 2021.
- [45] Maneet Singh, Shruti Nagpal, Richa Singh, and Mayank Vatsa. Dual directed capsule network for very low resolution image recognition. In *2019 IEEE/CVF International Conference on Computer Vision, ICCV 2019, Seoul, Korea (South), October 27 - November 2, 2019*, pages 340–349. IEEE, 2019. doi: 10.1109/ICCV.2019.00043. URL <https://doi.org/10.1109/ICCV.2019.00043>.
- [46] Wan-Chi Siu and Kwok-Wai Hung. Review of image interpolation and super-resolution. In *Asia-Pacific Signal and Information Processing Association Annual Summit and Conference, APSIPA 2012, Hollywood, CA, USA, December 3-6, 2012*, pages 1–10. IEEE, 2012. URL <https://ieeexplore.ieee.org/document/6411957/>.
- [47] Yi Sun, Yuheng Chen, Xiaogang Wang, and Xiaoou Tang. Deep learning face representation by joint identification-verification. In Zoubin Ghahramani, Max Welling, Corinna Cortes, Neil D. Lawrence, and Kilian Q. Weinberger, editors, *Advances in Neural Information Processing Systems 27: Annual Conference on Neural Information Processing Systems 2014, December 8-13 2014, Montreal, Quebec, Canada*, pages 1988–1996, 2014. URL <https://proceedings.neurips.cc/paper/2014/hash/e5e63da79fcd2bebbd7cb8bf1c1d0274-Abstract.html>.
- [48] Yifan Sun, Changmao Cheng, Yuhan Zhang, Chi Zhang, Liang Zheng, Zhongdao Wang, and Yichen Wei. Circle loss: A unified perspective of pair similarity optimization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6398–6407, 2020.
- [49] Ying Tai, Jian Yang, and Xiaoming Liu. Image super-resolution via deep recursive residual network. In *2017 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2017, Honolulu, HI, USA, July 21-26, 2017*, pages 2790–2798. IEEE Computer Society, 2017. doi: 10.1109/CVPR.2017.298. URL <https://doi.org/10.1109/CVPR.2017.298>.
- [50] Yonglong Tian, Dilip Krishnan, and Phillip Isola. Contrastive representation distillation. *CoRR*, abs/1910.10699, 2019. URL <http://arxiv.org/abs/1910.10699>.

- [51] Antonio Torralba, Robert Fergus, and William T. Freeman. 80 million tiny images: A large data set for nonparametric object and scene recognition. *IEEE Trans. Pattern Anal. Mach. Intell.*, 30(11):1958–1970, 2008. doi: 10.1109/TPAMI.2008.128. URL <https://doi.org/10.1109/TPAMI.2008.128>.
- [52] Hugo Touvron, Andrea Vedaldi, Matthijs Douze, and Hervé Jégou. Fixing the train-test resolution discrepancy. In Hanna M. Wallach, Hugo Larochelle, Alina Beygelzimer, Florence d’Alché-Buc, Emily B. Fox, and Roman Garnett, editors, *Advances in Neural Information Processing Systems 32: Annual Conference on Neural Information Processing Systems 2019, NeurIPS 2019, December 8-14, 2019, Vancouver, BC, Canada*, pages 8250–8260, 2019. URL <https://proceedings.neurips.cc/paper/2019/hash/d03a857a23b5285736c4d55e0bb067c8-Abstract.html>.
- [53] Frederick Tung and Greg Mori. Similarity-preserving knowledge distillation. In *2019 IEEE/CVF International Conference on Computer Vision, ICCV 2019, Seoul, Korea (South), October 27 - November 2, 2019*, pages 1365–1374. IEEE, 2019. doi: 10.1109/ICCV.2019.00145. URL <https://doi.org/10.1109/ICCV.2019.00145>.
- [54] Chien-Yi Wang, Ya-Liang Chang, Shang-Ta Yang, Dong Chen, and Shang-Hong Lai. Unified representation learning for cross model compatibility. In *31st British Machine Vision Conference 2020, BMVC 2020, Virtual Event, UK, September 7-10, 2020*. BMVA Press, 2020. URL <https://www.bmvc2020-conference.com/assets/papers/0195.pdf>.
- [55] Feng Wang, Xiang Xiang, Jian Cheng, and Alan Loddon Yuille. Normface: L2 hypersphere embedding for face verification. In *Proceedings of the 25th ACM international conference on Multimedia*, pages 1041–1049, 2017.
- [56] 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.
- [57] Xiaobo Wang, Shifeng Zhang, Shuo Wang, Tianyu Fu, Hailin Shi, and Tao Mei. Misclassified vector guided softmax loss for face recognition. In *The Thirty-Fourth AAAI Conference on Artificial Intelligence, AAAI 2020, The Thirty-Second Innovative Applications of Artificial Intelligence Conference, IAAI 2020, The Tenth AAAI Symposium on Educational Advances in Artificial Intelligence, EAAI 2020, New York, NY, USA, February 7-12, 2020*, pages 12241–12248. AAAI Press, 2020. URL <https://ojs.aaai.org/index.php/AAAI/article/view/6906>.
- [58] Zhangyang Wang, Shiyu Chang, Yingzhen Yang, Ding Liu, and Thomas S. Huang. Studying very low resolution recognition using deep networks. *CoRR*, abs/1601.04153, 2016. URL <http://arxiv.org/abs/1601.04153>.
- [59] 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.
- [60] Y. Wong, C. Sanderson, S. Mau, and B. C. Lovell. Dynamic amelioration of resolution mismatches for local feature based identity inference. In *2010 20th International Conference on Pattern Recognition*, 2010.

- [61] Junho Yim, Donggyu Joo, Ji-Hoon Bae, and Junmo Kim. A gift from knowledge distillation: Fast optimization, network minimization and transfer learning. In *2017 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2017, Honolulu, HI, USA, July 21-26, 2017*, pages 7130–7138. IEEE Computer Society, 2017. doi: 10.1109/CVPR.2017.754. URL <https://doi.org/10.1109/CVPR.2017.754>.
- [62] Xi Yin, Ying Tai, Yuge Huang, and Xiaoming Liu. FAN: feature adaptation network for surveillance face recognition and normalization. In Hiroshi Ishikawa, Cheng-Lin Liu, Tomás Pajdla, and Jianbo Shi, editors, *Computer Vision - ACCV 2020 - 15th Asian Conference on Computer Vision, Kyoto, Japan, November 30 - December 4, 2020, Revised Selected Papers, Part II*, volume 12623 of *Lecture Notes in Computer Science*, pages 301–319. Springer, 2020. doi: 10.1007/978-3-030-69532-3_19. URL https://doi.org/10.1007/978-3-030-69532-3_19.
- [63] Sergey Zagoruyko and Nikos Komodakis. Wide residual networks. In Richard C. Wilson, Edwin R. Hancock, and William A. P. Smith, editors, *Proceedings of the British Machine Vision Conference 2016, BMVC 2016, York, UK, September 19-22, 2016*. BMVA Press, 2016. URL <http://www.bmva.org/bmvc/2016/papers/paper087/index.html>.
- [64] Juan Zha and Hongyang Chao. TCN: transferable coupled network for cross-resolution face recognition*. In *IEEE International Conference on Acoustics, Speech and Signal Processing, ICASSP 2019, Brighton, United Kingdom, May 12-17, 2019*, pages 3302–3306. IEEE, 2019. doi: 10.1109/ICASSP.2019.8682384. URL <https://doi.org/10.1109/ICASSP.2019.8682384>.
- [65] Tianyue Zheng and Weihong Deng. Cross-pose lfw: A database for studying cross-pose face recognition in unconstrained environments. *Beijing University of Posts and Telecommunications, Tech. Rep*, 5:7, 2018.
- [66] Tianyue Zheng, Weihong Deng, and Jiani Hu. Cross-age LFW: A database for studying cross-age face recognition in unconstrained environments. *CoRR*, abs/1708.08197, 2017. URL <http://arxiv.org/abs/1708.08197>.
- [67] Mingjian Zhu, Kai Han, Enhua Wu, Qiulin Zhang, Ying Nie, Zhenzhong Lan, and Yunhe Wang. Dynamic resolution network. In Marc’Aurelio Ranzato, Alina Beygelzimer, Yann N. Dauphin, Percy Liang, and Jennifer Wortman Vaughan, editors, *Advances in Neural Information Processing Systems 34: Annual Conference on Neural Information Processing Systems 2021, NeurIPS 2021, December 6-14, 2021, virtual*, pages 27319–27330, 2021. URL <https://proceedings.neurips.cc/paper/2021/hash/e56954b4f6347e897f954495eab16a88-Abstract.html>.
- [68] Shizhan Zhu, Sifei Liu, Chen Change Loy, and Xiaoou Tang. Deep cascaded bi-network for face hallucination. In Bastian Leibe, Jiri Matas, Nicu Sebe, and Max Welling, editors, *Computer Vision – ECCV 2016*, pages 614–630, Cham, 2016. Springer International Publishing.