LAE : Long-tailed Age Estimation

Zenghao Bao^{1,2,*}, Zichang Tan^{4,5,*}, Yu Zhu^{4,5}, Jun Wan^{1,2,†}, Xibo Ma^{1,2}, Zhen Lei^{1,2,3}, and Guodong Guo^{4,5}

¹ CBSR&NLPR, Institute of Automation, Chinese Academy of Sciences {baozenghao2020, jun.wan}@ia.ac.cn

² School of Artificial Intelligence, University of Chinese Academy of Sciences

{xibo.ma, zlei}@nlpr.ia.ac.cn

³ Centre for Artificial Intelligence and Robotics, Hong Kong Institute of Science Innovation, Chinese Academy of Sciences

⁴ Institute of Deep Learning, Baidu Research

⁵ National Engineering Laboratory for Deep Learning Technology and Application {tanzichang,zhuyu05,guoguodong01}@baidu.com

Abstract. Facial age estimation is an important yet very challenging problem in computer vision. To improve the performance of facial age estimation, we first formulate a simple standard baseline and build a much strong one by collecting the tricks in pre-training, data augmentation, model architecture, and so on. Compared with the standard baseline, the proposed one significantly decreases the estimation errors. Moreover, long-tailed recognition has been an important topic in facial age datasets, where the samples often lack on the elderly and children. To train a balanced age estimator, we propose a two-stage training method named Long-tailed Age Estimation (LAE), which decouples the learning procedure into representation learning and classification. The effectiveness of our approach has been demonstrated on the dataset provided by organizers of Guess The Age Contest 2021.

Keywords: Deep Learning \cdot Age Estimation \cdot Long-tailed Recognition.

1 Introduction

As one of the most concerning topics in facial analysis, age estimation remains challenging for the uncontrolled nature of the aging process, a significant variance among faces in the same age range, and a high dependency of aging traits on a person. Since age estimation from face images is an inevitable problem in several real applications, enhancing age estimation performance has also become an urgent problem to be solved. In recent years, deep learning methods [10, 18, 5, 19, 21], which automatically learn the effective image representations, have gained many successes in facial age estimation.

^{*}Co-First Author

[†] Corresponding Author

However, the most promising methods [10, 18, 19, 2, 12, 11] use ensembles of DCNNs, making the obtained classifier not usable in real applications. These methods usually require prohibitive computational resources, complex training procedures caused by the plurality of neural networks, and substantial training sets, which are not simply collectible.

Motivated by this concern, we build a strong baseline based on a single neural network by collecting some efficient tricks in pre-training, model architecture, data augmentation, and so on. The strong baseline achieves impressive performance compared to the standard baseline. Nevertheless, the resulting estimator is still not a balance age estimator and usually shows poor performance in classes with fewer samples.

Unfortunately, existing methods for age estimation [10, 5, 21] still lack in longtailed scenarios. The long-tail distribution of the visual world poses great challenges for deep learning-based classification models on how to handle the data imbalance problem. Inspired by the existing work [9], we decouple the learning procedure into representation learning and classification to train a balanced estimator.

In the representation learning stage, we obtain a robust feature extractor to distinguish different ages. Furthermore, in the classification stage, we freeze the model except for the FC layer and train the model on balanced MIVIA [6] dataset obtained by class-balanced sampling [15]. Moreover, we innovatively introduce Mean Square Error (MSE) loss into this stage for further narrowing down the standard deviation of different age groups. By applying this, we obtain a model that performs well at all ages.

To summarize, the main contributions are listed as follows:

- We collect some effective training tricks to build a strong baseline based on a single neural network, which effectively enhances age estimation performance.
- We propose a novel LAE method to deal with the age estimation in longtailed scenarios. The resulting estimator performs well at all ages.
- Extensive experiments on MIVIA indicate the effectiveness of LAE.

2 Proposed Method

2.1 Overview

In our implementation, we break the neural network into two components, a feature extractor θ and a classifier ψ . The classifier consists of a linear layer and a softmax layer. Before explaining the proposed method in detail, we first define the key symbols. Assume we have an age dataset D. The *i*th sample in the D is denoted by (x_i, y_i, z_i) , where x_i denote the input image, y_i denote the age label, and z_i denote the label distribution. The label distribution of the sample is based on a typical Gaussian distribution, *i.e.*, $z_i^k = \frac{1}{\sqrt{2\pi\sigma}} exp(-\frac{(k-y_i)^2}{2\sigma^2})$ where $k \in [0, ..., 100]$ and standard deviation σ is set to 1.

2.2 Standard Baseline

In this subsection, we provide a standard baseline used for comparison. During the training stage, the pipeline includes the following steps:

1. We initialize the ResNet18 [8] with pre-trained parameters on ImageNet [4] and change the dimension of the fully connected layer to K. And K is set to 101 in our experiments.

2. We resize each image into 224 \times 224 pixels and each image is flipped horizontally with 0.5 probability.

3. Each image is decoded into 32-bit floating point raw pixel values in [0, 1]. Then we normalize RGB channels by subtracting 0.485, 0.456, 0.406 and dividing by 0.229, 0.224, 0.225, respectively.

4. Following the work [5], the network is trained with a KL divergence and ℓ_1 loss. KL divergence is calculated by measuring the distance between label distribution z_i and predicted age distribution \hat{z}_i , and ℓ_1 loss aims to narrow the distance between the predicted age \hat{y}_i and the ground truth label y_i . Specifically, the predicted age \hat{y}_i is obtained by using an expectation regression on the softmax outputs.

5. SGD is adopted to optimize the model. The initial learning rate is set to be 0.005 and is decreased by 0.1 at the 20th, 40th, and 60th epoch. Totally there are 75 epochs.

2.3 Training Tricks

In this subsection, we introduce some effective tricks for training. After expanding all the tricks, the standard baseline becomes a strong baseline.

Model Architecture : Compared to the previous models, EfficientNetV2 [17] has a faster training speed and better parameter efficiency. In our submission, we use the EfficientNetV2-M as the backbone of the feature extractor.

Pre-training: Pre-training has a significant impact on the performance of an age estimator. Before training the network on age datasets, we pretrain the model on MS-celeb-1M [7], with only using identity labels. Following the work [20], which AM-softmax loss is employed as the loss function. In this way, a large number of face images are employed to enhance the network's discriminating ability with obtaining a good initialization.

Batch Size: The batch size also shows its impact on performance during training. Compared to 16, 64, and 128, keep 32 samples in one GPU shows the best performance. So the default batch size for us is 256 since we use 8 GPUs for training.

Onecycle Scheduler: Existing methods [18, 5] for age estimation usually require many epochs to convergence, which undoubtedly brings a more significant time cost and occupation of computing resources. The existence of Onecycle [16] can make networks be trained much faster than before. In our case, it reduces the number of training epochs from the original 75 to 24.

Data Augmentation: As a data-space solution to the problem of limited data, data augmentation is capable of enhancing the size and quality of training datasets to avoid overfitting and make the resulting model more generalization. In our case, we applied data augmentation as follows: first, a random patch of the images is selected and resized to 224×224 with a random horizontal flip, followed by a color jitter, consisting of a random sequence of brightness, contrast, saturation, hue adjustments. Finally, RandAugment [3] is applied, and we set N=2 and M=9, where N denotes the number of transformations to apply, and M denotes the magnitude of the applied transformations.

Fig. 1. The distribution of the dataset samples over the age is depicted on the left. The long-tailed distribution of the MIVIA is depicted on the right, sorted by the number of samples in class. Head-class denotes high-frequency class, and tail-class denotes low-frequency class. In practice, the number of samples per class generally decreases from head to tail classes.



2.4 Long-tailed Age Estimation

The long-tail distribution of the visual world poses great challenges for deep learning-based classification models on how to handle the data imbalance problem [9, 13]. As the result, the resulting model usually shows poor performance in classes with fewer samples. This phenomenon is reflected in most of the age datasets. As shown in Fig. 2, the age distribution of the MIVIA is imbalanced, in line with the premise of long-tailed recognition. Inspired by related research in the field of long-tailed recognition [9], we decouple the learning procedure into representation learning and classification to train a relatively balanced classifier, which performs well at all ages.

Representation Learning Stage: In this stage, we follow the pipeline of strong baseline for training a robust feature extractor to distinguish different ages. To this end, we finetune the model on the MIVIA training set. By feeding the feature maps into the classifier, we obtain the prediction distribution \hat{z}_i^k , which indicates the probability of classifying x_i to age k. Then, we use KL divergence to measure the label distribution loss L_{ld} between ground-truth label distribution \hat{z}_i^k [5]:

$$L_{ld} = \sum_{k=1}^{K} z_i^k \log \frac{z_i^k}{\hat{z}_i^k} \tag{1}$$

The predicted age \hat{y}_i is obtained by using a expectation regression[14] on the output distribution. It can be denoted as $\hat{y}_i = \sum_{k=1}^{K} k \hat{z}_i^k$. Moreover, we employ a loss L_{er} to further narrowing the distance between the predicted age label and the ground-truth label :

$$L_{er} = |y_i - \sum_{k=1}^{K} k \hat{z}_i^k|$$
(2)

where $|\cdot|$ denotes ℓ_1 loss.

Finally, we integrate all the above loss functions and train them simultaneously. Referring to the losses in Eq. 1, 2 as L_{ld}, L_{er} , respectively, then overall we optimize the feature extractor with L_{fe} :

$$L_{fe} = L_{ld} + \lambda L_{er} \tag{3}$$

where λ denotes the weight of L_{er} .

Classification Stage: To retrain the classifier without changing the parameters of the feature extractor, we freeze the model except for the FC layer and train on balanced MIVIA obtained by class-balanced sampling, which means each class has an equal probability of being selected. One can see this as a two-stage sampling strategy, where a class is first selected uniformly from the set of classes, and then an instance from that class is subsequently uniformly sampled [15].

Moreover, we innovatively introduce MSE loss into this stage for further narrowing down the standard deviation of different age groups. To this end, we optimize the classifier with L_{cl} :

$$L_{cl} = L_{ld} + (L_{er} - mae)^2$$
(4)

where *mae* denotes the mean absolute error over the validation set.

$$mae = \sum_{i=1}^{M} \frac{|y_i - \hat{y}_i|}{M}$$
(5)

where M is the total number of validating images, y_i and \hat{y}_i denote the groundtruth age and the predicted age, respectively. It is worth noting that the *mae* here results from the representation learning stage on the validation set.

3 Experiments

In this section, we first introduce the datasets and the evaluation metrics from our experiments. Then we provide implementation details for our LAE method, describe experimental setups, and discuss results.

3.1 Dataset

The MIVIA Age Dataset is composed of 575,073 images of more than 9,000 identities. Those images have been extracted from the VGGFace2 [1] dataset and annotated with age by means of a knowledge distillation technique, making the dataset very heterogeneous in terms of face size, illumination conditions, facial pose, gender, and ethnicity.

3.2 Metric

An index called Age Accuracy and Regularity (AAR) is introduced for taking into account accuracy and regularity:

$$AAR = max(0;7 - MAE) + max(0;3 - \sigma)$$
(6)

where MAE denotes the mean absolute error on the entire test set, and σ is obtained by:

$$\sigma = \sqrt{\frac{\sum_{j=1}^{8} (MAE^{j} - MAE)^{2}}{8}}$$
(7)

where MAE^{j} denotes the mae and is computed over the samples whose real age is in age group j. The details of age group are shown as below.

Table 1. The details of the eight age groups.

MAE^1	MAE^2	MAE^3	MAE^4	MAE^5	MAE^{6}	MAE^7	MAE^8
1-10	11-20	21-30	31-40	41-50	51-60	61-70	71-81

3.3 Implementation Details

We use an SGD optimizer with a MultiStepLR scheduler over 18 epochs in the pre-training stage, with the milestones of [7, 14, 16]. Moreover, we use an SGD optimizer with a OneCycle schedule over 24 epochs and 8 epochs in the representation and classification stage. We set the base learning rate to 0.005 and 0.001, scaled linearly with the batch size (LearningRate = base × BatchSize/256). In addition, we use a global weight decay parameter of 5e-4. The default batch size is set to 256. According to the work [22], the images are aligned with five landmarks (including two eyes, nose tip, and two mouth corners). For a fair comparison, we fix the $\lambda = 1$ during all experiments.



Fig. 2. Examples of apparent age estimation using LAE on MIVIA testing images.

3.4 Experimental Results

Since the test set is not available during the competition, we randomly divide the MIVIA into a training set and a validation set at a ratio of 4:1 to validate the effectiveness of the proposed methods. The result is shown as below. The standard baseline achieves 1.97 on MAE. Then, we add model architecture, pretraining, and data augmentation, respectively. The final strong baseline boosts more performance than other models. Finally, the bag of tricks makes the estimator achieves 1.71 on MAE and 7.18 on AAR. The results (such as in Fig. 2) shows that our proposed solution is able to predict the age of faces as well as people can or even better.

Moreover, the strong baseline is also the result of our representation learning stage. The estimator performs well at head-classes but shows poor performance at tail-classes. As a result, it also performs poorly on the σ and the AAR. Then, we add the classification stage into the training procedure, and the final estimator shows the best performance on tail-classes, σ , and AAR on the MIVIA.

The results on test sets are provided by the organizers. And it is worth noting that the submission is trained on the whole MIVIA.

Table 2. The performance of different models is evaluated on MIVIA. Baseline-S stands for the standard baseline introduced in Section 2.2. MAE^1 and MAE^8 stand for age groups consisting of tail-classes, and MAE^3 stands for head-classes. * denotes the model is evaluated on the test set, and other models are all evaluated on the validation set. The best results on the validation set are in bold.

Model	MAE	MAE^1	MAE^3	MAE^8	σ	AAR
Baseline-S	1.97	7.16	1.61	3.45	1.93	6.10
+model architecture	1.93	6.47	1.59	2.99	1.65	6.42
+pre-training	1.78	5.09	1.45	2.73	1.33	6.89
+data augmentation	1.85	5.53	1.52	2.99	1.37	6.78
Representation (Strong Baseline)	1.71	4.66	1.40	2.59	1.11	7.18
Classification	1.89	2.69	1.62	2.11	0.37	7.74
Submission*	1.86	/	/	/	0.20	7.94

4 Conclusions

In this paper, we first formulate a standard baseline and build a strong baseline by collecting the tricks in pre-training, model architecture, data augmentation, and so on. The bag of tricks significantly enhances the age estimation performance. The resulting estimator achieves the best performance on head-classes. Second, we propose an LAE method, which decouples the learning procedure into representation learning and classification, to obtain a more balanced estimator. The proposed method performs well at all ages, significantly outperforming the human reference.

Acknowledgments

This work was supported by the Chinese National Natural Science Foundation Projects #61961160704, #61876179, the External cooperation key project of Chinese Academy Sciences # 173211KYSB20200002, the Key Project of the General Logistics Department Grant No.AWS17J001, Science and Technology Development Fund of Macau (No. 0010/2019/AFJ, 0008/2019/A1 0025/2019/A-KP0019/2018/ASC).

References

- Cao, Q., Shen, L., Xie, W., Parkhi, O.M., Zisserman, A.: Vggface2: A dataset for recognising faces across pose and age. In: 2018 13th IEEE international conference on automatic face & gesture recognition (FG 2018). pp. 67–74. IEEE (2018)
- Carletti, V., Greco, A., Percannella, G., Vento, M.: Age from faces in the deep learning revolution. IEEE transactions on pattern analysis and machine intelligence 42(9), 2113–2132 (2019)
- 3. Cubuk, E.D., Zoph, B., Shlens, J., Le, Q.V.: Randaugment: Practical automated data augmentation with a reduced search space. In: CVPRW (2020)
- Deng, J., Dong, W., Socher, R., Li, L.J., Li, K., Fei-Fei, L.: Imagenet: A largescale hierarchical image database. In: 2009 IEEE conference on computer vision and pattern recognition. pp. 248–255. Ieee (2009)
- Gao, B.B., Zhou, H.Y., Wu, J., Geng, X.: Age estimation using expectation of label distribution learning. In: IJCAI. pp. 712–718 (2018)
- Greco, A., Saggese, A., Vento, M., Vigilante, V.: Effective training of convolutional neural networks for age estimation based on knowledge distillation. Neural Computing and Applications pp. 1–16 (2021)
- Guo, Y., Zhang, L., Hu, Y., He, X., Gao, J.: Ms-celeb-1m: A dataset and benchmark for large-scale face recognition. In: European conference on computer vision. pp. 87–102. Springer (2016)
- He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: CVPR. pp. 770–778 (2016)
- Kang, B., Xie, S., Rohrbach, M., Yan, Z., Gordo, A., Feng, J., Kalantidis, Y.: Decoupling representation and classifier for long-tailed recognition. arXiv preprint arXiv:1910.09217 (2019)

- 10. Niu, Z., Zhou, M., Wang, L., Gao, X., Hua, G.: Ordinal regression with multiple output cnn for age estimation. In: CVPR (2016)
- Othmani, A., Taleb, A.R., Abdelkawy, H., Hadid, A.: Age estimation from faces using deep learning: A comparative analysis. Computer Vision and Image Understanding 196, 102961 (2020)
- Punyani, P., Gupta, R., Kumar, A.: Neural networks for facial age estimation: a survey on recent advances. Artificial Intelligence Review 53(5), 3299–3347 (2020)
- Ren, J., Yu, C., Sheng, S., Ma, X., Zhao, H., Yi, S., Li, H.: Balanced meta-softmax for long-tailed visual recognition. arXiv preprint arXiv:2007.10740 (2020)
- Rothe, R., Timofte, R., Van Gool, L.: Dex: Deep expectation of apparent age from a single image. In: Proceedings of the IEEE international conference on computer vision workshops. pp. 10–15 (2015)
- Shen, L., Lin, Z., Huang, Q.: Relay backpropagation for effective learning of deep convolutional neural networks. In: European conference on computer vision. pp. 467–482. Springer (2016)
- Smith, L.N., Topin, N.: Super-convergence: Very fast training of neural networks using large learning rates. In: Artificial Intelligence and Machine Learning for Multi-Domain Operations Applications. vol. 11006, p. 1100612. International Society for Optics and Photonics (2019)
- 17. Tan, M., Le, Q.V.: Efficientnetv2: Smaller models and faster training. arXiv preprint arXiv:2104.00298 (2021)
- Tan, Z., Wan, J., Lei, Z., Zhi, R., Guo, G., Li, S.Z.: Efficient group-n encoding and decoding for facial age estimation. IEEE TPAMI (2018)
- Tan, Z., Yang, Y., Wan, J., Guo, G., Li, S.Z.: Deeply-learned hybrid representations for facial age estimation. In: IJCAI. pp. 3548–3554 (2019)
- Wang, H., Wang, Y., Zhou, Z., Ji, X., Gong, D., Zhou, J., Li, Z., Liu, W.: Cosface: Large margin cosine loss for deep face recognition. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 5265–5274 (2018)
- Zeng, X., Huang, J., Ding, C.: Soft-ranking label encoding for robust facial age estimation. IEEE Access (2020)
- Zhang, K., Zhang, Z., Li, Z., Qiao, Y.: Joint face detection and alignment using multitask cascaded convolutional networks. IEEE Signal Processing Letters (2016)