Quantifying Facial Age by Posterior of Age Comparisons

Yunxuan Zhang¹
 zhangyunxuan@sensetime.com
 Li Liu¹
 liuli@sensetime.com

Cheng Li¹
 chengli@sensetime.com

Chen Change Loy²
 ccloy@ie.cuhk.edu.hk

¹ SenseTime Group Limited

² Department of Information Engineering,
 The Chinese University of Hong Kong

Abstract

We introduce a novel approach for annotating large quantity of in-the-wild facial images with high-quality posterior age distribution as labels. Each posterior provides a probability distribution of estimated ages for a face. Our approach is motivated by observations that it is easier to distinguish who is the older of two people than to determine the person's actual age. Given a reference database with samples of known ages and a dataset to label, we can transfer reliable annotations from the former to the latter via human-in-the-loop comparisons. We show an effective way to transform such comparisons to posterior via fully-connected and SoftMax layers, so as to permit end-to-end training in a deep network. Thanks to the efficient and effective annotation approach, we collect a new large-scale facial age dataset, dubbed 'MegaAge', which consists of 50,000 images. With the dataset, we train a network that jointly performs ordinal hyperplane classification and posterior distribution learning. Our approach achieves state-of-the-art results on popular benchmarks such as MORPH2, Adience, and the newly proposed MegaAge.

1 Introduction

Is there a moment when you try to guess the age of your new friend or a stranger and you seriously doubt your estimation? You may try very hard to find clues on his/her face – determining the number of wrinkles surround the eyes or sagging skin above the upper eyelids. Even with multiple trials, your guess sometimes may still far from the true answer.

The aforementioned difficulty is frequently encountered during the preparation of facial age datasets, which are needed for training models for automatic facial age estimation. Given face samples whose age is unknown, it is notoriously hard and unreliable to label them with actual age. The best we can do is resort to providing approximate age range annotation, which is adopted by Adience dataset [1], one of the largest in-the-wild facial age databases. An alternative is to collect photos from subjects whose age is known. Under such restrictions, the sample size is usually small and the images are constrained by the environment from which we capture them. It is thus inevitably hard to generalise a model learned from such a

Figure 1: (a) Comparing the proposed target annotation (also the prediction) with existing schemes. In contrast to current schemes that either provide an actual age label or age range, we introduce a new annotation that faithfully captures the distribution of estimated ages. The annotation can be obtained from age comparisons between query and reference images. (b) We propose a network that can jointly estimate ordinal classifications and age posterior distribution as outputs.

dataset to novel scenes. Early databases, such as MORPH [17] and FG-NET [16], fall into this category.

While it is hard to guess the actual age of an individual, one may find it effortless to tell apart who is elder or younger when two persons stand next to each other. Age comparison is usually deemed easier since we can use the appearance of a person as a reference while examining the other. Nevertheless, it remains an open problem on how we can exploit such comparisons as ground-truths for training a robust and accurate age estimation model. In this paper, we introduce an effective way of exploiting age comparisons for labelling massive quantity of in-the-wild face images. We further propose a new paradigm for mapping a series of such comparisons to meaningful age distribution posterior that can be used to train an age estimation deep network in an end-to-end manner. We detail our contributions as follows.

1) New large-scale age posterior dataset - As the first contribution of this study, we propose a novel age annotation approach given face images collected in the wild. Specifically, we first conduct a user study and show that humans are poor in guessing the actual age of a face but they perform much better at comparing the age of two faces. The observation then leads to our new approach – given an image with unknown age, we compare the image with some selected images from a reference database with known age labels. We treat each comparison as an independent random event and determine the age distribution as the posterior given multiple observed events. The posterior is set as the target label of the image. As shown in Fig. 1(a), the posterior faithfully reflects the true age with an estimated range to capture uncertainty, which otherwise cannot be achieved by existing annotation schemes. The proposed approach enables us to extend reliable age annotations from any existing datasets (e.g., FG-NET [16]), which are captured in constrained environments, to a massive in-the-wild dataset with human-in-the-loop. The resulting dataset, called *MegaAge*, consists of 50,000 faces annotated with age posterior distributions. The images are randomly selected from MegaFace [18].

2) Joint cost sensitive and posterior losses - The second contribution of this work is a new deep convolutional network architecture that learns from age posterior distribution. The network also generates age posterior as its output. The network is inspired by existing ordinal hyperplane methods [9, 12, 15], which learn multiple binary ‘is this face older than age k ?’ classifiers for each possible age k . Ordinal hyperplane method has shown powerful capa-

bility of keeping ordinal information of age labels. Extending from the idea, our network consists of a ordinal hyperplane module that captures ordinal information. Different from ordinal hyperplane methods, our network further maps the responses of the module to generate age posterior distribution. We show that the mapping can be accomplished by using a fully-connected layer and a softmax layer. The proposed network is unique in that it can be jointly trained with two losses, namely, the cost sensitive loss that is typically applied for learning an ordinal hyperplanes ranker [10], and the newly proposed Kullback-Leibler (KL) divergence loss that supervises the learning of posterior distribution. The two losses are complementary in nature. The cost sensitive loss enforces the learning of ordinal relationships, while the KL divergence loss ensures the estimated age to fall within a distribution and captures uncertainties. In addition, with the KL divergence loss, our network can generate smoother ordinal classifications (as a side product) in comparison to conventional ordinal hyperplane methods, as shown in Fig. 1(b).

Extensive experiments are conducted on MORPH2, Adience and MegaAge datasets. Thanks to the novel annotation approach, we collected a large quantity of training samples with good quality to train the proposed network. Our approach outperforms state-of-the-art approaches by reducing over 12% and 10% exact error on the MORPH2 [13] and Adience [13] datasets, respectively¹. We also set a new benchmark on the proposed MegaAge.

2 Quantifying Age Posterior from Age Comparisons

In this section, we first present a user study and then discuss how we could employ the observations of the user study and present our approach for mapping age comparisons to age posterior distribution.

2.1 How Good can Human Predict Facial Age?

Test I - Guessing Actual Age: In the first experiment, we asked 30 volunteers to guess the actual age given face images randomly drawn from the FG-NET [14] dataset. All participants were not familiar with the dataset. A total of 1002 images were presented to each of the participants. Figure 2(a) shows the results. The participants were fairly poor at estimating the actual age through face appearance on this dataset. The average recall rate within an ± 3 years old is only 43.2%.

Test II - Comparing Age: In the second experiments, the same group of participants were asked to compare the age of two faces from FG-NET. Figure 2(b) shows the relationship between age difference (face A minus face B), and the probability that participants find A is older than B. It is observed that when the age difference is over 10 years, participants gained over 95% accuracy in predicting the relative age order. Even when the age difference is only 5 years, participants still performed well with an accuracy of 85%.

Discussion: Albeit the performances of the two tests are not directly comparable due to their different natures, the study does reflect that humans are more comfortable in telling apart persons of different ages, especially when the age difference is large. Another important observation is that human’s performance on age comparisons can be approximated by a *logistic function*, as exhibited by the curve depicted in Fig. 2(b).

¹On MORPH, we achieved mean absolute error of 2.87, compared to 3.27 reported in [13]. On Adience, we achieved 56.01% exact accuracy compared to 50.70% reported in [13].

Figure 2: Results of the user study: (a) Test I – volunteers were asked to guess the actual age given an image. (b) Test II – volunteers were asked to compare the age of two persons.

2.2 From Age Comparisons to Age Posterior

We learn from the user study presented in Sec. 2.1 that each pair of age comparison can be approximated by a logistic function. Specifically, we treat each comparison between a query face I and a reference face I_{ref} of known age k as a random event denoted as $C_k \in \{0, 1\}$, where $C_k = 1$ indicates I is older than I_{ref} . The likelihood of event C_k given that face I has an age a is denoted $P(C_k|a)$,

$$P(C_k|a) = \begin{cases} \sigma(\beta(a-k))/Z & \text{if } C_k = 1 \\ \sigma(\beta(k-a))/Z & \text{if } C_k = 0 \end{cases}, \quad (1)$$

where $\sigma(\cdot)$ is a logistic function with $\sigma(x) = 1/(1 + \exp(-x))$, Z is a partition function with $Z = P(C_k = 0|a) + P(C_k = 1|a)$. The parameter β is a value that can be obtained by fitting the curve generated from our user study (Fig. 2(b)). We use $\beta = 0.36$ throughout our experiments. Figure 3 depicts a few examples that show the likelihood $P(C_k|a)$ derived from age comparisons across with different reference images. Let's denote the events of M comparisons as $C_{k_1}, C_{k_2}, \dots, C_{k_M}$, we can compute the age posterior probability as

$$P(a|\{C_{k_m}\}) \propto P(a) \prod_{m=1}^M P(C_{k_m}|a), \quad (2)$$

where $P(a)$ is the prior of age, which we assume to be uniform in this study. An example of posterior distribution is shown in Fig. 3. It is worth pointing out that a narrower posterior distribution can be obtained if more comparisons are conducted.

Discussion: One may expect different logistic curves for different age ranges, since the difficulty level may increase from comparing younger ages to older ages. In this paper, we simplified the assumption on the logistic function by using a single β , but with satisfactory results. Future work can explore different logistic functions across ages.

2.3 MegaAge Dataset

In this section we detail the construction of the *MegaAge* dataset with posterior distribution serving as the target label of each image. We randomly sampled query images from the challenging MegaFace dataset [18]², which contains a million unconstrained photos that capture more than 690K different individuals. We adopted the widely used FG-NET as our reference database. The labelling process takes the same procedures as illustrated in Fig. 3. Specifically, given a query image sampled from MegaFace, we compared it with six face images selected from the reference database. We introduced a few constraints in the

²<http://megaface.cs.washington.edu/>

Figure 3: An illustration to show the age posterior distribution $P(a|\{C_{k_m}\})$ as a product of a prior $P(a)$ and all likelihoods $P(C_{k_m}|a)$.

Figure 4: Examples of images in the proposed MegaAge dataset and the associated age posterior label. Each blue vertical line represents a comparison (between the query face and a reference face) and its direction of supporting the posterior. The last figure shows an outlier that arises due to disagreement of annotators.

selection process to ensure meaningful comparisons. Firstly, the reference images need to have the same gender as the query image so as to avoid gender bias. An accurate gender classifier with an accuracy rate of over 97% was used. Secondly, we roughly estimated the age of the query image using an existing model trained on MORPH and selected three images smaller than the estimated age and another three that are larger than the estimated age. Selection is random apart from the constraints. Age posterior was generated following the method presented in Sec. 2.2. We repeated this process to label all query images with posterior. Outliers with a posterior distribution with 90% confidence interval larger than 15 years old were discarded. This kind of outliers constituted only 3% of queries. More than 70% of queries have their posterior’s confidence interval less than 8 years old. The final number of images we collected with posterior ground-truth is 50,000. Note that actual age ground-truth can be easily derived from the posterior by taking the mode of the distribution. Some of the samples are shown in Fig. 4.

3 Deep Learning from Age Posterior

We now describe the deep convolutional network that learns from age posterior. Figure 5 illustrates the proposed network. The network comprises a truncated VGG CNN network for extracting 128-dimensional features from a given image. Further details on this CNN is

Figure 5: The proposed network can be trained end-to-end. It is supervised by two signals: cost sensitive loss and KL divergence loss. The former is used to regularise the ordinal hyperplane module. The latter provides further constraints to ensure the estimated age to fall within the desired age distribution.

given in Sec. 4. Apart from the feature extraction layer, the network consists of an ordinal hyperplane module and an age posterior distribution module. The network is jointly trained with cost sensitive loss and KL divergence loss. We detail the modules as follows.

3.1 Ordinal Hyperplane Module

Ordinal hyperplane method [14, 15] has been shown very powerful in capturing ordinal information of age labels. To gain advantage from this nice characteristic, we formulate an ordinal hyperplane module in our network. The module can be trained with cost sensitive loss, and the gradient can be back-propagated to further update the parameters of the feature extraction CNN. It is noted that our network can be trained without using the cost sensitive loss (using KL divergence loss alone), but we observed better performance when the network is trained with the loss as a regulariser.

Design: The module comprises of a fully-connected (FC) layer, an output layer, and a sigmoid layer, as illustrated in Fig. 5. The output layer treats the age labels y_i as a rank order, $y_i \in \{1, \dots, K\}$, where K is the number of labels (typically set as 70 in our approach). The FC layer basically learns K classification function $f_k(x)$, to model the confidence of “face’s age is larger than k ”, that is $f_k(I) = \langle \mathbf{w}_k, \phi(I) \rangle$, where $\phi(I)$ is the feature of a face image, and $\mathbf{w}_k$ is the linear weight of the classifier encapsulated in the FC. The sigmoid layer converts the ordinal classifications to responses as an input for the subsequent age posterior distribution module, which is described next in Sec. 3.2.

Loss: We follow [14, 15] to use cost sensitive loss as an auxiliary loss to regularise our network. The cost sensitive loss is a relaxed regression loss, which can be written as

$$L_{\text{hyper}} = \sum_k \text{Cost}_k(a_{gt}) \|f_k(I) - \mathbf{1}[a_{gt} > k]\|_2^2, \quad (3)$$

where a_{gt} is the ground-truth age, $\mathbf{1}[\cdot]$ is an indicator function, which equals 1 for true argument, and 0 otherwise. Here $\text{Cost}_k(a_{gt})$ is a truncated cost [14] which equals 0 when $|a_{gt} - k| < L$. Here we set L as 3.

3.2 Age Posterior Distribution Module

Note that the ordinal hyperplane module only captures the ordinal information of age labels but falls short of representing the uncertainty of an estimated range of ages. Here we show that it is possible to extend the ordinal hyperplane module to a probabilistic one.

Design: Recall that we learn classifiers $f_k(I)$ in the ordinal hyperplane module. The response of a k -th classifier can be treated as the probability of a random event $E_k \in \{0, 1\}$, where $E_k = 1$ when the classifier thinks that a face is older than age k , that is $f_k(I) = P(E_k = 1|I) = \sigma(\mathbf{w}_k^T \phi(I))$, where $\mathbf{w}_k$ is a learnable weight and $\phi(I)$ is the deep convolutional feature. Next, we assume all K classifiers are independent, and estimate the age posterior given all $\{E_k\}_{k=1}^K$,

$$\begin{aligned} P(a|I) &= P(a|\{E_k\}) = \frac{1}{Z} P(a) \prod_{k=1}^K P(E_k|a) \\ &= \frac{1}{Z} P(a) \prod_{k=1}^K P(E_k = 1|a)^{f_k(I)} P(E_k = 0|a)^{(1-f_k(I))}. \end{aligned} \quad (4)$$

Notice here we approximate $P(E_k|a)$ as an exponential combination $P(E_k = 1|a)^{f_k(I)} \cdot P(E_k = 0|a)^{(1-f_k(I))}$ instead of linear one $f_k(I)P(E_k = 1|a) + (1 - f_k(I))P(E_k = 0|a)$. If we take log on both sides of Eq. (4), the equation turns to

$$\log P(a|I) = -\log Z + \log P(a) + \sum_k (f_k(I) \log P(E_k = 1|a) + (1 - f_k(I)) \log P(E_k = 0|a)). \quad (5)$$

Similar to the assumption we made in Sec. 2.2, we assume that a likelihood can be modelled as a logistic function, *e.g.*, $P(E_k = 1|a) = \sigma(\beta(k - a))$, which only depends on the age difference. Here we use the same β coefficient as in Sec. 2.2. Note that the posterior $\log P(a|I)$ is just a linear function of ordinal hyperplane classifier $f_k(I)$, which means we can conveniently compute $P(a|I)$ by linking a fully connected layer after the ordinal hyperplane module, and followed by a SoftMax layer for normalisation. This design is parameter free and easy to implement.

Loss: We can train the network end-to-end using the ground-truth age posterior distribution, denoted as $P_{gt}(a)$. We employ KL-divergence between two distribution as our loss,

$$L_{KL} = D_{KL}(P_{gt}(a)||P(a|I)) = -\sum_a P_{gt}(a) \log P(a|I) - \text{Const}. \quad (6)$$

Here $P(a|I)$ is the soft output of our network. There are three possible ways to prepare for $P_{gt}(a)$: (1) for dataset that is labelled only with actual age, *e.g.*, MORPH, we set P_{gt} as a sharp Gaussian distribution with $\sigma = 2$. (2) For category based dataset, *e.g.*, Adience, the confidence interval of each distribution of each sample is equivalent to the age range each specific age category covers. (3) for MegaAge dataset, we use the ground-truth age posterior.

4 Experiments

Datasets: We evaluate our method on MORPH2 [14], Adience [8] and MegaAge.

1) *MORPH2* contains more than 55,000 face images of 13,000 individuals aged from 16 to 77 years. On average, each individual has more than 3 images and exact age was given. MORPH2 is the largest publicly available aging dataset. We follow the experimental setting in [14], where the data is randomly divided into 80%/20% exclusive training/test partitions.

2) *Adience* provides 26,580 images, which cover an age range from 0 to 70. The samples are divided into 8 age groups (0-2, 4-6, 8-13, 15-20, 25-32, 38-43, 48-53, 60-). We follow the standard protocol [8] to perform a 5-fold cross validation.

3) *MegaAge*, as described in Sec. 2.3, contains 50,000 images encompassing ages from 0 to 70. We reserve 10,000 images as test data. Each sample was labelled with age posterior as the ground-truth.

Table 1: We first train a deep model with the training partition of MORPH2 and Adience (indicated by ‘0%’). We then gradually add MegaAge training samples to the initial training pool to boost the model’s performance.

Percentage of MegaAge Added	0%	50%	100%
Adience (Exact)	53.43 $\pm$ 3.83	53.92 $\pm$ 4.78	54.52 $\pm$ 3.74
Adience (1-off)	88.18 $\pm$ 3.26	91.29 $\pm$ 2.64	91.29 $\pm$ 1.63
MegaAge (CA(3))	39.70	52.27	63.09
MegaAge (CA(5))	53.26	76.48	80.79

Evaluation Metric: The performance on MORPH2 is evaluated by mean average error (MAE). As for Adience, we report the exact mean accuracy on 8-class classification. We also provide the 1-off accuracy [10], *i.e.*, a prediction is considered correct if it hits the exact age group or its neighbouring groups. For MegaAge, we employ cumulative accuracy as our metric, that is, $CA(n) = K_n/K \times 100$, where K_n is the number of test images whose absolute estimated error is smaller than n . To align with the results of Adience, of which the range of each age group is around 5, we report $CA(3)$ and $CA(5)$ in our experiments.

Network Architecture: In this study we use a truncated VGG-network, which only has one forth of filter number in each convolutional layer of the original VGG network [24]. Thus the computational cost is reduced to 1/16 of the original. Our network runs at 10 frames per second (FPS) on a i7 desktop and 50 FPS on a GTX 970 GPU. We initialise all networks with a face verification model trained on MS-Celeb-1M dataset [8], of which the result on LFW [10] is 99.3%.

4.1 Evaluating The Quality of MegaAge Annotations

To show that the proposed annotation approach (Sec. 2.2) is capable of generating high-quality ground-truth to boost the performance of existing models, we conduct an experiment in which we gradually add more MegaAge data to the initial pool of training set. The initial pool is formed by the training partition of MORPH2 and Adience. Here we use a variant of the proposed network – we use the truncated VGG network and train it with cost sensitive loss only. Prediction is made via ordinal classification and the aggregation rule presented by [10]. We feed this network with increasing MegaFace training data from 0% to 100%. To report results on Adience, the method first estimates the exact age and then assigns the estimation to the nearest age group defined. It is evident from Table 1 that annotations of MegaAge is of high quality ones and they can improve the deep learning based method with a considerable margin.

4.2 Comparisons with Existing Methods and Ablation Study

Results on Adience and MegaAge. In this experiment we compared our method with state-of-the-art methods [4, 9, 13] on both the Adience and MegaAge datasets. For the Cumulative Attribute method [4], we retrained the method using VGG face verification features trained on MS-Celeb-1M. For Deep CNN [13], we reported the results given in the original paper. In addition to the aforementioned comparisons, we also evaluated the importance of different losses used in the training of our network:

1) w/o L_{KL} - we removed the KL divergence loss and used only the cost sensitive loss. In this setting, the age estimation is given by aggregating the ordinal classifications of the ordinal

Table 2: A comparison with existing methods and ablation study. The Cumulative Attribute [10] method was re-trained with VGG face verification features.

Test Data Metric	Adience		MegaAge	
	Exact	1-off	CA(3)	CA(5)
Existing methods:				
LBP+FLBP+Dropout 0.8 [10]	45.10 ± 2.60	79.50 ± 1.40	—	—
Deep CNN [13]	50.70 ± 5.10	84.70 ± 2.20	—	—
Cumulative Attribute [10]	52.34 ± 3.72	89.34 ± 1.79	64.21	81.44
Proposed method and variants:				
w/o L_{hyper}	49.46 ± 5.51	82.99 ± 3.74	59.44	77.19
w/o L_{KL}	54.52 ± 3.74	91.29 ± 1.63	63.09	80.79
Full model	56.01 ± 4.41	91.42 ± 2.15	64.51	82.31

Figure 6: a) We illustrate the predictions of our full model on the Adience dataset. The box plot shows the median, first/third quartiles, and min/max on actual age predictions, and how they fall into different age groups defined by Adience. (b) Some qualitative results are shown. Underscored images are failure cases (under the 1-off evaluation metric), which are caused by heavy occlusion and wrong faces detected.

hyperplane module following [10]. Note that one will need actual age labels for using the cost sensitive loss. To enable us for training our network with MegaAge, we used MAP estimate, the mode of the ground-truth distribution, as an approximation to the actual age ground-truth.

2) w/o L_{hyper} - we discarded the cost sensitive loss and used only the KL-divergence loss.

3) Full model - our full method with both the L_{KL} and L_{hyper} .

Results in Table 2 show that the performance of our network drops considerably without either of the cost sensitive loss, L_{Hyper} , or the KL divergence loss, L_{KL} , suggesting the complementary nature and importance of these two losses. As shown earlier in Fig. 1(b), the KL divergence loss actually helps the ordinal hyperplane module to produce smoother classifications. The proposed full method outperforms state-of-the-art Deep CNN [13] and Cumulative Attribute (re-trained with deep features) approaches [10]. Qualitative results on Adience are depicted in Fig. 6.

Results on MORPH2. Here we provide additional results on MORPH2, despite the performance on it were long saturated. We use the same setting following [13]. The following baselines are tested: 1) *DEX w/o IMDB-WIKI* [13]: A VGG-19 Net pretrained on ImageNet 2) *DEX w/ IMDB-WIKI* [13]: The same model as 1) but the ImageNet pre-trained model is further pretrained on a large-scale age dataset, named as IMDB-WIKI, which contains

500,000 celebrities available on IMDB and Wiki. 3) *Ours w/o IMDB-WIKI*: our full model with mixed loss, i.e., cost sensitive loss and KL divergence loss. The model is pre-trained on a face verification task. 4) *Ours w/ IMDB-WIKI*: a VGG-19 network which is the same as 2) but it employs the proposed mixed loss. The model was pre-trained on IMDB-WIKI.

Table 3: Comparative results based on mean absolute error (MAE). Results are migrated from [18]. The proposed method achieves the state-of-the-art performance on MORPH2 (* indicates different data split).

Method	MORPH2[18]	Method	MORPH2[18]
Human workers[9]	6.30	DIF[9]	3.80
AGES[9]	8.83	MTWGP[24]	6.28
CA-SVR[9]	5.88	SVR[9]	5.77
OHRank[9]	5.69	DLA[18]	4.77
[18]	4.25*	[9]	4.18*
[9]	3.92*	[23]	3.63*
[18]	3.45	[18]	3.27
DEX w/o IMDB-WIKI	3.25	DEX w/ IMDB-WIKI	2.68
Ours w/o IMDB-WIKI	2.87	Ours w/ IMDB-WIKI	2.52

Table 3 summarises our result on MORPH2. With the additional supervision using the age posterior distribution by KL-divergence loss, we reduced MAE on MORPH2 dataset by 0.16, when both DEX and our model are trained with IMDB-WIKI. The improvement is significant considering that MORPH2 is a rather mature benchmark with saturated performance.

5 Conclusion

We have presented a novel annotation scheme for facial age database. Instead of labelling actual age or age range as practiced in existing studies, we label each face with age posterior distribution to better capture the uncertainty and ordinal information of age labels. The distribution can be conveniently and reliably obtained through a small number of comparisons between a query face with reference faces from another database. Experimental results suggest the good quality of the annotations. It is noteworthy that the proposed annotation scheme is not limited to facial age labelling, it is suitable for other subjective labelling targets, e.g., degree of continuous facial expression change. In this study, we also introduced a way to train a deep network directly using age posterior ground-truths via the KL divergence loss. The loss is found to be complementary with the commonly applied cost sensitive loss.

Acknowledgement: This work is supported by SenseTime Group Limited and the General Research Fund sponsored by the Research Grants Council of the Hong Kong SAR (CUHK 14224316).

References

- [1] Kuang-Yu Chang, Chu-Song Chen, and Yi-Ping Hung. Ordinal hyperplanes ranker with cost sensitivities for age estimation. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 585–592, 2011.

- [2] Ke Chen, Shaogang Gong, Tao Xiang, and Chen Change Loy. Cumulative attribute space for age and crowd density estimation. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 2467–2474, 2013.
- [3] Eran Eidinger, Roei Enbar, and Tal Hassner. Age and gender estimation of unfiltered faces. *IEEE Transactions on Information Forensics and Security*, 9(12):2170–2179, 2014.
- [4] Xin Geng, Zhi-Hua Zhou, and Kate Smith-Miles. Automatic age estimation based on facial aging patterns. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 29(12):2234–2240, 2007.
- [5] Guodong Guo and Guowang Mu. Simultaneous dimensionality reduction and human age estimation via kernel partial least squares regression. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 657–664, 2011.
- [6] Guodong Guo and Guowang Mu. A framework for joint estimation of age, gender and ethnicity on a large database. *Image and Vision Computing*, 32(10):761–770, 2014.
- [7] Guodong Guo, Yun Fu, Charles R Dyer, and Thomas S Huang. Image-based human age estimation by manifold learning and locally adjusted robust regression. *IEEE Transactions on Image Processing*, 17(7):1178–1188, 2008.
- [8] 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*, 2016.
- [9] Hu Han, Charles Otto, Xiaoming Liu, and Anil K Jain. Demographic estimation from face images: Human vs. machine performance. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 37(6):1148–1161, 2015.
- [10] Gary B Huang, Manu Ramesh, Tamara Berg, and Erik Learned-Miller. Labeled faces in the wild: A database for studying face recognition in unconstrained environments. Technical report, Technical Report 07-49, University of Massachusetts, Amherst, 2007.
- [11] Ivan Huerta, Carles Fernández, and Andrea Prati. Facial age estimation through the fusion of texture and local appearance descriptors. In *European Conference on Computer Vision Workshop*, pages 667–681, 2014.
- [12] Ira Kemelmacher-Shlizerman, Steven M Seitz, Daniel Miller, and Evan Brossard. The MegaFace benchmark: 1 million faces for recognition at scale. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2016.
- [13] Gil Levi and Tal Hassner. Age and gender classification using convolutional neural networks. In *IEEE Conference on Computer Vision and Pattern Recognition Workshop*, 2015.
- [14] Changsheng Li, Qingshan Liu, Jing Liu, and Hanqing Lu. Learning ordinal discriminative features for age estimation. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2012.

- [15] Zhenxing Niu, Mo Zhou, Le Wang, Xinbo Gao, and Gang Hua. Ordinal regression with multiple output CNN for age estimation. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2016.
- [16] Gabriel Panis and Andreas Lanitis. An overview of research activities in facial age estimation using the FG-NET aging database. In *European Conference on Computer Vision*, 2014.
- [17] Karl Ricanek and Tamirat Tesafaye. MORPH: A longitudinal image database of normal adult age-progression. In *IEEE International Conference on Automatic Face & Gesture Recognition*, 2006.
- [18] Rasmus Rothe, Radu Timofte, and Luc Van Gool. Deep expectation of real and apparent age from a single image without facial landmarks. *International Journal of Computer Vision*, pages 1–14, 2016.
- [19] Rasmus Rothe, Radu Timofte, and Luc Van Gool. Some like it hot-visual guidance for preference prediction. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 5553–5561, 2016.
- [20] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*, 2014.
- [21] Xiaolong Wang, Rui Guo, and Chandra Kambhampettu. Deeply-learned feature for age estimation. In *IEEE Winter Conference on Applications of Computer Vision*, pages 534–541, 2015.
- [22] Huei-Fang Yang, Bo-Yao Lin, Kuang-Yu Chang, and Chu-Song Chen. Automatic age estimation from face images via deep ranking. In *British Machine Vision Conference*, 2013.
- [23] Dong Yi, Zhen Lei, and Stan Z Li. Age estimation by multi-scale convolutional network. In *Asian Conference on Computer Vision*, pages 144–158, 2014.
- [24] Yu Zhang and Dit-Yan Yeung. Multi-task warped gaussian process for personalized age estimation. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 2622–2629, 2010.