

An Effective Deep Network for Head Pose Estimation without Keypoints

Chien Thai^a, Viet Tran, Minh Bui, Huong Ninh and Hai Tran

Computer Vision Department, Optoelectronics Center, Viettel Aerospace Institute, Vietnam

Keywords: Head Pose Estimation, Knowledge Distillation, Convolutional Neural Network.

Abstract: Human head pose estimation is an essential problem in facial analysis in recent years that has a lot of computer vision applications such as gaze estimation, virtual reality, driver assistance. Because of the importance of the head pose estimation problem, it is necessary to design a compact model to resolve this task in order to reduce the computational cost when deploying on facial analysis-based applications such as large camera surveillance systems, AI cameras while maintaining accuracy. In this work, we propose a lightweight model that effectively addresses the head pose estimation problem. Our approach has two main steps. 1) We first train many teacher models on the synthesis dataset - 300W-LPA to get the head pose pseudo labels. 2) We design an architecture with the ResNet18 backbone and train our proposed model with the ensemble of these pseudo labels via the knowledge distillation process. To evaluate the effectiveness of our model, we use AFLW-2000 and BIWI - two real-world head pose datasets. Experimental results show that our proposed model significantly improves the accuracy in comparison with the state-of-the-art head pose estimation methods. Furthermore, our model has the real-time speed of ~300 FPS when inferring on Tesla V100.

1 INTRODUCTION


Head pose estimation (HPE) is an important problem in facial analysis that has been extensively researched in recent years. Its application can be widely observed in lots of intelligent computer vision systems including virtual reality (Kumar et al., 2017), driver assistance (Schwarz et al., 2017; Murphy-Chutorian et al., 2007), gaze estimation (Murphy-Chutorian and Trivedi, 2008), human-computer interaction (Seemann et al., 2004; Wang et al., 2019) and smart city surveillance.

The objective of head pose estimation is to accurately identify the orientation of heads of individuals found in images. Existing methods to solve this problem can be divided into two primary categories: landmark-based approaches (Cao et al., 2014; Lathuilière et al., 2017; Fanelli et al., 2011; Xiong and De la Torre, 2015; Sun et al., 2013; Xin et al., 2021; Bulat and Tzimiropoulos, 2017; DeMenthon and Davis, 1995) and landmark-free approach (Ruiz et al., 2018; Yang et al., 2019; Zhou and Gregson, 2020; Chang et al., 2017). Landmark-based methods use facial keypoints extracted by landmark detectors to regress the head pose angle. Recently, these approaches have achieved remarkable results since the

usage of deep neural networks has greatly enhanced the quality of landmark detectors. However, the problem remains challenging due to the fact that not only a minor error of landmark detectors may adversely affect the head pose estimation but learning the relation between the geometric distribution of facial landmarks and head poses is not a trivial task. Furthermore, using landmark detection as a preprocessing step imposes a computational burden for the whole process of estimating head angle which hinders its usage for real-time applications. Landmark-free methods, on the other hand, directly predict the head poses from images without detecting facial keypoints which results in their fast execution time.

In addition to the above approaches, some works utilize depth information from depth cameras (Meyer et al., 2015; Fanelli et al., 2011; Mukherjee and Robertson, 2015; Martin et al., 2014). Although this approach provides a prominent result, it still has some limitations. The depth cameras are sensitive to illumination change and light conditions so that they often yield substandard results in an uncontrolled environment. Moreover, they are very expensive and use more storage and transfer time, so they are often impractical for real-time applications.

Because of the importance of the head pose estimation problem and in order to minimize the pro-

^a  <https://orcid.org/0000-0002-5098-6862>

cessing time of the model when deploying on large systems or embedded platforms, our goal is to design a lightweight architecture that solves this task while still guaranteeing remarkable performance. For having a compact and simple model, our network uses ResNet18 architecture as a backbone. The contributions of our work can be summarized as follows:

- We address a major mistake found in HopeNet (Ruiz et al., 2018) in which annotated face boxes are mislabeled. We prove that correcting those mislabeled boxes can significantly improve the accuracy of the head pose estimation task.
- An end-to-end deep architecture designed to solve head pose estimation problem is proposed. A lightweight model is trained to this task via the knowledge distillation process.
- Experiments are conducted to evaluate the performance of our method on two challenging head pose datasets (BIWI and AFLW-2000). Our method achieves state-of-the-art performance when evaluating on the head pose dataset.

The rest of the paper is organized as follows: Section 2 puts forward some related works on head pose estimation problem. In section 3, we present our proposed method. Section 4 discusses the datasets, experiments, results, and ablation study. Finally, the conclusion and future work are discussed in Section 5.

2 RELATED WORKS

Convolutional Neural Networks. (CNNs) are widely used in computer vision tasks and gradually replace the traditional image processing methods. CNN is designed to automatically learn the spatial features of the image by using convolution kernels. With many convolutional layers, deep networks can extract high-level semantic features. He et al. (He et al., 2016) propose the Residual Network to train the much deeper convolutional neural network. ResNet uses a skip connection between the current layer and the previous layer which can learn the identity mapping and solve the vanishing gradient problem. Because of its powerful and simple architecture, ResNet and its variants (Xie et al., 2017; Zhang et al., 2020; Gao et al., 2019) are widely used in many computer vision applications and deliver high performance.

Human Head Pose Estimation. has been researched over the past 25 years with many different approaches. Appearance Template (Niyogi and Freeman, 1996; Beymer, 1994; Sherrah et al., 2001; Ng

and Gong, 2002; Sherrah et al., 1999) is the method that compares the input image with a set of labeled templates and assigns it to the most similar template. Detector arrays (Huang et al., 1998; Zhang et al., 2006; Jones and Viola, 2003) estimate head pose by training multiple face detectors for the different discrete poses.

Many approaches are based on facial landmarks from the input image to estimate the head pose. With the progress of landmarks detection, landmark-based methods demonstrate superior performance. Dementhon et al. (DeMenthon and Davis, 1995) proposed Pose from Orthography and Scaling with Iterations which determines the head pose by 3D computer vision techniques for the given 2D face landmarks. FAN (Bulat and Tzimiropoulos, 2017) using deep neural network to estimate 3D face models. EVA-GCN (Xin et al., 2021) constructs a landmark-connection graph and leverages the Graph Convolution Network (Yan et al., 2018) to learn the nonlinear relationships between head poses and distribution of facial keypoints.

Multi-task methods combine the head pose estimation problem with other related facial analysis problems, such as face detection, keypoints detection. Some works show that learning with related tasks yields better results than learning individual tasks independently (Chen et al., 2014; Kumar et al., 2017; Zhu and Ramanan, 2012; Ranjan et al., 2017b). KEPLER (Kumar et al., 2017) predicts face detection and pose estimation jointly by using Heatmap-CNN to capture structured global and local features. Hyperface (Ranjan et al., 2017a) presents a convolutional neural network for simultaneous face detection, landmarks localization, pose estimation, and gender recognition.

Gu et al. (Gu et al., 2017) proposed a dynamic facial analysis that uses a recurrent neural network. They improve head pose estimation and facial landmarks localization by leveraging the time dimension from videos instead of a single frame.

For accurate head pose estimation, some methods utilize 3D information of depth images. Meyer et al. (Meyer et al., 2015) perform head pose estimation by registering 3D morphable models to depth images, using the particle swarm optimization and the iterative closest point algorithm. Fanelli et al. (Fanelli et al., 2011) using Random Regression Forests to regress the head pose estimation of depth images.

Recent works directly predict the Euler angles from a single RGB image by using a deep neural network and achieve prominent performance. HopeNet (Ruiz et al., 2018) proposed a multi-loss framework that combines binned pose classification and regres-




Figure 1: The overview of the head pose model. The original image is passed through the face detector to get the bounding box of the objective face. The detected face is padded to a squared image and resized to 112x112. The head pose model extract 62 dimensions distribution vector for the given image. The predicted pose is calculated by the expectation of this vector. For each Euler angle, the classification loss is the cross-entropy loss between distribution vector and one-hot vector, the regression loss is the mean square error of ground truth and predicted pose.

sion loss for each Euler angle. By using a very stable softmax layer and cross-entropy for binned classification loss, the network obtained robust neighborhood prediction of the head pose. FSA-Net (Yang et al., 2019) employs the soft stagewise regression scheme by training classification and regression objectives of the features from multiple stages. It provides a compact model and accurate prediction. WHENet (Zhou and Gregson, 2020) proposed wrapped loss to estimate the full 360-degree range of yaw angle. Our proposed network has similar architecture to HopeNet (Ruiz et al., 2018), but has a smaller model size and achieves better performance on two challenging head pose datasets - BIWI and AFLW-2000.

3 PROPOSED METHOD

In this section, we describe the major disadvantage of previous work and the method to mitigate this problem. After that, we explain the proposed method to construct an effective head pose estimation model via knowledge distillation process.

The head pose estimation problem can be mathematically formulated as: Given a set of training images $X = \{x^i | i = 1..N\}$ and ground truth $Y = \{y^i | i = 1..N\}$, where N is number of images, and y^i is 3D vector of image x^i corresponding to three Euler angles (yaw, pitch, roll), the goal is to find a function F so that the absolute difference between $F(x)$ and the real head pose y for the given image x as small as possible.

Inspired by HopeNet (Ruiz et al., 2018), we design a network using a multi-loss framework to solve this problem. HopeNet casts the regression problem of head pose estimation as a classification problem by dividing the poses range into 66 bins, each bin contains 3 units of degree. The predicted pose is the expected value of classes distribution.

Further investigating HopeNet, we found that it is the preprocessing data that hinders its performance. They loosely crop around the bounding box of a face on the image and resize the cropped image to 224x224

before fitting the model. Because the height of the face bounding box is often longer than the width, it slightly changes the real head pose and causes a negative effect on the training and testing phase.

To mitigate this problem, we padded the bounding box of face to squared shape. Given a bounding box (x_1, y_1, x_2, y_2) , the padding size k is calculated by $|x_2 - x_1 - y_2 + y_1|$ (the absolute difference between width and height). If the height $h = x_2 - x_1$ is longer than the width $w = y_2 - y_1$, the new coordinates of bounding box (x'_1, y'_1, x'_2, y'_2) are:

$$x'_1 = x_1$$

$$x'_2 = x_2$$

$$y'_1 = y_1 - [k/2]$$

$$y'_2 = y_2 + [k/2]$$

and vice versa. After getting the square image of faces, we resize it to (112, 112) in order to decrease the computation cost when training and inference.

Unlike HopeNet, we divide the poses range from -93 to 93 into 62 bins for each Euler angle. The classification loss of angle is cross-entropy loss between softmax output of model and pose's corresponding one-hot vector:

$$\mathcal{L}_{cls}^{angle} = \sum_{i=1}^N y^i * \log(\hat{y}^i) \quad (1)$$

where y^i and \hat{y}^i are respectively one-hot vector of pose and predicted softmax output for given input x^i .

The predicted pose of x^i is expected values of softmax output that is denoted by r^i . The regression loss of angle is mean squared error between the ground truth labels y^i and the predicted pose r^i :

$$\mathcal{L}_{reg}^{angle} = \frac{1}{N} \sum_{i=1}^N \|r^i - y^i\|^2 \quad (2)$$

The total loss is composed by three separate losses, each loss is calculated by the sum of classification and regression loss of angle, as following:

$$\mathcal{L} = \sum \mathcal{L}_{cls}^{angle} + \mathcal{L}_{reg}^{angle} \quad (3)$$

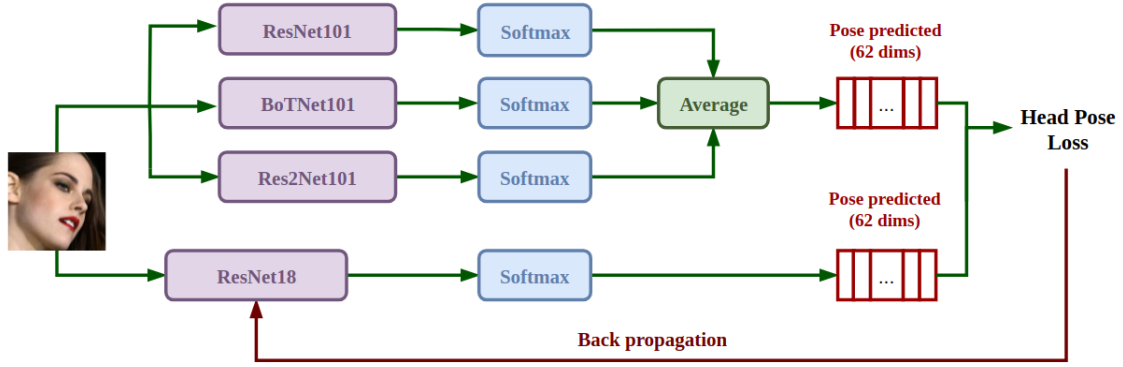


Figure 2: The overview of proposed method. The student model using ResNet18 backbone. The head pose loss is the sum of Kullback-Leibler Divergence loss between softmax output of student model and ensemble output of head pose teacher models on each of yaw, pitch, roll angle. The total loss is sum of distillation loss of three Euler angles.

where $angle \in \{yaw, pitch, roll\}$

The above method uses hard labels to train head pose estimation models. Inspired by (Hinton et al., 2015), we use knowledge distillation to construct a compact model while enhancing the performance of this task. Our network uses ResNet18 (He et al., 2016) as the backbone, a simple and small architecture which is trained to match the output of head pose teacher models (pseudo label). With supervised learning, models are trained to match the same labels but with the different initiation and architectures, they will focus on distinctive features. So, we ensemble outputs of several strong head pose models to get more informative teacher features.

Given $N_{teacher}$ head pose models, we ensemble by calculating mean regression outputs of them. It is equal to the expected value of mean softmax outputs of these models. So, the output after ensemble n teacher head pose models is:

$$y_i^{ens} = \frac{1}{N_{teacher}} \sum_{j=1}^{N_{teacher}} \hat{y}_i^j \quad (4)$$

where \hat{y}_i^j is softmax output of head pose teacher model j for given image x^i

The loss function for head pose task is Kullback-Leibler Divergence between softmax output of student model \hat{y}^t and output ensemble of n teacher models y^{ens} :

$$\mathcal{L}_{headpose} = - \sum_{i=1}^N y_i^{ens} * \log\left(\frac{\hat{y}_i^t}{y_i^{ens}}\right) \quad (5)$$

Because head pose estimation is a challenging task, we found that the stronger model with a lot of parameters and computation cost, the more model's capacity to achieve good results. Base on the performance on ImageNet dataset (Deng et al., 2009), we train three head pose teacher models from

scratch whose backbones are chosen respectively as ResNet101 (He et al., 2016), BotNet101 (Srinivas et al., 2021), and Res2Net101 (Gao et al., 2019). After that, we train a head pose model with backbone ResNet18 by the aforementioned head pose knowledge distillation strategy.

In our experiment, we observed that the big models (teacher models) often give larger probabilities to the bins in the proximity of the truth bin and smaller scores to the ones far away. This is valuable information (i.e. the faces in a bin are more likely the faces in its neighbor bins) but it has very little effect on the cross entropy cost function during training if the probabilities are so close to zero. This means the soft targets of the teacher models attain a variety of information than one-hot labels, which helps the small model (student model) learn easily. So, we argue that the distilled model can preserve the generalization of the teacher models and reaches highly accurate results.

4 EXPERIMENTAL RESULTS

In this section, we describe the datasets for training and testing, implementation, results, comparisons with other state-of-the-art methods and the ablation study.

4.1 Dataset

Headpose Dataset: In our experiment, we use three popular datasets for the head pose estimation problem: 300W-LPA (Hsu et al., 2019), AFLW-2000 (Zhu and Ramanan, 2012), and BIWI (Fanelli et al., 2011) datasets. 300W-LPA is a synthetically expanded dataset that provides over 350000 images across large poses. The AFLW-2000 dataset provides head pose

Table 1: Mean average error of Euler angles across both state-of-the-art landmark-based and landmark-free methods on the BIWI and AFLW2000 dataset.

Model	BIWI				AFLW-2000			
	Yaw	Pitch	Roll	MAE	Yaw	Pitch	Roll	MAE
KEPLER (Kumar et al., 2017)	8.80	17.3	16.2	13.9	-	-	-	-
FAN (Bulat and Tzimiropoulos, 2017)	8.53	7.48	7.63	7.89	6.36	12.3	8.71	9.12
Dlib (Kazemi and Sullivan, 2014)	16.8	13.8	6.19	12.2	23.1	13.6	10.5	15.8
3DDFA (Zhu et al., 2016)	-	-	-	-	5.40	8.53	8.25	7.39
EVA-GCN (Xin et al., 2021)	4.01	4.78	2.98	3.92	4.46	5.34	4.11	4.64
HopeNet ($\alpha = 2$) (Ruiz et al., 2018)	5.17	6.98	3.39	5.18	6.47	6.56	5.44	6.16
HopeNet ($\alpha = 1$) (Ruiz et al., 2018)	4.81	6.61	3.27	4.90	6.92	6.64	5.67	6.41
SSR-Net-MD (Yang et al., 2018)	4.49	6.31	3.61	4.65	5.14	7.09	5.89	6.01
FSA-Caps-Fusion (Yang et al., 2019)	4.27	4.96	2.76	4.00	4.50	6.08	4.64	5.07
WHENet-V (Zhou and Gregson, 2020)	3.60	4.10	2.73	3.48	4.44	5.75	4.31	4.83
EHPNet (Ours)	3.68	4.03	2.57	3.43	3.23	5.54	3.88	4.15




Figure 3: Some examples of face image from the datasets. The first row is from the 300W-LPA (Hsu et al., 2019) which is a synthetically dataset. The second row and third row are respectively from the AFLW-2000 (Zhu et al., 2016) and BIWI (Fanelli et al., 2011) - two real-world datasets.

ground truth and corresponds to 68 landmark points among 2000 3-D face images. Images in the AFLW-2000 dataset have large pose annotation and various lighting conditions. BIWI dataset uses a Kinect v2 device to record RGB-D video of different subjects. It contains 24 videos of 20 subjects across different head poses. There are roughly 15000 samples in this dataset, each sample contains RGB and depth images, and pose annotations were created by using depth information.

Following HopeNet, we use 300-LPA for training while testing on AFLW-2000 and BIWI - two real-world datasets. In our case, we only use RGB images for training in these datasets. We run RetinaFace

(Deng et al., 2019) on all images to get the coordinate of the bounding box of faces.

4.2 Implementation

For better estimating on low-resolution face images, we augment the head pose training dataset by random downsampling and upsampling to original image size, randomly adjust brightness and contrast, blur by random Gaussian kernel. We randomly flip the image and relabel the yaw and roll angle of the flip image to $-yaw$ and $-roll$ to get more training data.

We use Pytorch for implementing the proposed network. We use 100 epochs to train the teacher networks with hard labels and 200 epochs for the knowledge distillation process. The chosen optimizer is Adam with initial learning-rate is $1e-4$. The learning rate is reduced by the cosine annealing strategy. The experiments are performed on a computer with a Tesla V100 GPU.

4.3 Results

We compare our proposed network with other state-of-the-art head pose estimation methods on BIWI and AFLW datasets. KEPLER (Kumar et al., 2017), FAN (Bulat and Tzimiropoulos, 2017), Dlib (Kazemi and Sullivan, 2014) and EVA-GCN (Xin et al., 2021) are landmark-based methods. KEPLER (Kumar et al., 2017) uses a modified GoogleNet to predict facial landmark points and pose at the same



Figure 4: The scatter plot between yaw, pitch, roll values and errors on AFLW-2000 dataset.




Figure 5: The scatter plot between yaw, pitch, roll values and errors on BIWI Dataset.

time. Dlib (Kazemi and Sullivan, 2014) is a face library that uses an ensemble of regression trees to detect landmarks. FAN (Bulat and Tzimiropoulos, 2017) is a state-of-the-art landmark detection method. EVA-GCN (Xin et al., 2021) is a state-of-the-art landmark-based method which constructs a landmark-connection graph and leverages the Graph Convolution Network (Yan et al., 2018) to learn the nonlinear relationships between head poses and distribution of facial keypoints. HopeNet (Ruiz et al., 2018), FSANet (Yang et al., 2019) and WHENet (Zhou and Gregson, 2020) are landmark free methods which treat the regression problem as classification problem by dividing the poses range to different classes. The α coefficient of HopeNet is the weight of the regression losses.

Table 1 shows the comparisons of our proposed network with these above models respectively on BIWI and AFLW-2000 datasets. The evaluation metric is the mean absolute error.

As shown in Table 1, our proposed EHPNet achieves state-of-the-arts on both AFLW-2000 and BIWI datasets. It outperforms the previous state-of-the-art WHENet by respectively 14.9% and 1.15%. EHPNet has a similar architecture with WHENet and HopeNet but uses a smaller model and image size. Even so, it has more significant improvement compared to WHENet and HopeNet. Furthermore, our model has a real-time speed of 300FPS when inferring on Tesla V100.

Figure 6 shows some visualization of predicted

images on BIWI, AFLW-2000. Besides achieving good results on images with variant poses and various lighting conditions, the network can also predict well on face images with high occlusion.




Figure 6: Results of the proposed network. The blue line indicates the front of the face, the green line pointing downward direction and the red line pointing to the side. The first row is the prediction on the BIWI dataset. The second row is the estimation result on images with various lighting conditions of AFLW-2000.

The scatter diagrams in Figure 4 and Figure 5 show the influence of the pose’s value on prediction results for each angle on the AFLW-2000 and BIWI datasets. On AFLW-2000, the yaw angle has a smaller mean absolute error than pitch and roll angles and has a stable prediction for pose range. For the pitch and roll angles, the model tends to predict well if the value of pose is as close to 0. As shown in Figure 4, there is some prediction that has a very big error. We find that this happens because the head poses on AFLW-2000 are provided by using 3D landmarks, so some examples can have a big difference in a pose if viewed as RGB images. On the BIWI dataset, the discrepancy

Table 2: The impact of different backbone and distillation training on head pose estimation models. The evaluation metric is the mean absolute error of Euler angles.

Backbone	BIWI				AFLW-2000			
	Yaw	Pitch	Roll	MAE	Yaw	Pitch	Roll	MAE
ResNet18	3.969	4.849	2.869	3.897	3.785	5.642	4.238	4.555
ResNet101	3.680	3.945	2.755	3.460	3.249	5.276	3.821	4.115
BotNet101	3.876	4.066	2.528	3.489	3.559	5.109	3.697	4.135
Res2Net101	3.827	3.939	2.669	3.478	3.223	5.080	3.556	3.953
Ensemble	3.688	3.859	2.508	3.352	3.169	5.009	3.560	3.913
Distilled ResNet18	3.683	4.033	2.571	3.429	3.226	5.345	3.876	4.148

between ground truth and prediction is not significant but the trending is slightly changed. The model predicts better on yaw and roll angle. The higher pose's value of pitch angle leads to face occlusion in the image and makes the model confused. As shown in Figure 5, many samples from the BIWI dataset don't follow the trend. For example, a sample which has a yaw angle value close to zero has a maximum error. In our experiments, we observed that, although it has a small value of the yaw angle, it has a large pitch and roll so the face can be occluded, and leads to wrong predictions.

4.4 Ablation Study

We have conducted the ablation study when changing the backbone and using a pseudo label from teacher models for the head pose estimation task. As shown in Table 2, the result has significant improvement when using padding instead of resizing the cropped face image like HopeNet. By using ensemble, the mean absolute errors are slightly decreased on both BIWI and AFLW datasets. The small head pose model achieves better accuracy when training via the knowledge distillation process. With ResNet18 as a backbone, the pose model using pseudo labels from the output of ensemble many teacher models is better than the same model using the hard label. The distilled head pose model has equivalent results to its teacher, even better.

In our experiment, we observed that these models can predict the same result for the yaw angle. But for pitch and roll angle, the complex head pose model works better. Among three teacher models, each one can predict better at a specific pose interval (i.e. model ResNet101 achieves the smallest error of yaw with 3.680, while Res2Net101 outperforms the

two others with the pitch error of 3.939 and the last one BotNet101 attains the best roll error of 2.528), but worse at the others. By preserving the generalization of each model, the ensemble results have a stable prediction on the poses range. Overall, training the baseline model with hard targets leads to severe overfitting, whereas training the same model with ensemble soft targets is able to recover better generalization and achieves competitive results.

5 CONCLUSIONS

In this paper, we have presented an EHPNet, which can directly, accurately and robustly predict the head rotation from a single RGB image. This is achieved by mitigating the disadvantages of previous work, along with distilling the knowledge from many robust head pose estimation teacher models. By using ResNet18 architecture as a backbone, the model is compact and usable in many computer vision applications. The proposed network outperforms both landmark-based and landmark-free methods and achieves state-of-the-art results on both the AFLW-2000 and BIWI datasets, with higher respectively 14.9% and 1.15% than WHENet.

In the future, we would like to use a lower computation network as well as reduce the input image resolution. Besides, more effective knowledge distillation techniques will be used to help the student model achieve better accuracy.

REFERENCES

- Beymer, D. (1994). Face recognition under varying pose. In *CVPR*, volume 94, page 137. Citeseer.

- Bulat, A. and Tzimiropoulos, G. (2017). How far are we from solving the 2d & 3d face alignment problem?(and a dataset of 230,000 3d facial landmarks). In *Proceedings of the IEEE International Conference on Computer Vision*, pages 1021–1030.
- Cao, X., Wei, Y., Wen, F., and Sun, J. (2014). Face alignment by explicit shape regression. *International journal of computer vision*, 107(2):177–190.
- Chang, F.-J., Tuan Tran, A., Hassner, T., Masi, I., Nevatia, R., and Medioni, G. (2017). Faceposenet: Making a case for landmark-free face alignment. In *Proceedings of the IEEE International Conference on Computer Vision Workshops*, pages 1599–1608.
- Chen, D., Ren, S., Wei, Y., Cao, X., and Sun, J. (2014). Joint cascade face detection and alignment. In *European conference on computer vision*, pages 109–122. Springer.
- DeMenthon, D. F. and Davis, L. S. (1995). Model-based object pose in 25 lines of code. *International journal of computer vision*, 15(1-2):123–141.
- Deng, J., Dong, W., Socher, R., Li, L.-J., Li, K., and Fei-Fei, L. (2009). Imagenet: A large-scale hierarchical image database. In *2009 IEEE conference on computer vision and pattern recognition*, pages 248–255. Ieee.
- Deng, J., Guo, J., Zhou, Y., Yu, J., Kotsia, I., and Zafeiriou, S. (2019). Retinaface: Single-stage dense face localization in the wild. *arXiv preprint arXiv:1905.00641*.
- Fanelli, G., Weise, T., Gall, J., and Van Gool, L. (2011). Real time head pose estimation from consumer depth cameras. In *Joint pattern recognition symposium*, pages 101–110. Springer.
- Gao, S., Cheng, M.-M., Zhao, K., Zhang, X.-Y., Yang, M.-H., and Torr, P. H. (2019). Res2net: A new multi-scale backbone architecture. *IEEE transactions on pattern analysis and machine intelligence*.
- Gu, J., Yang, X., De Mello, S., and Kautz, J. (2017). Dynamic facial analysis: From bayesian filtering to recurrent neural network. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1548–1557.
- He, K., Zhang, X., Ren, S., and Sun, J. (2016). Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778.
- Hinton, G., Vinyals, O., and Dean, J. (2015). Distilling the knowledge in a neural network. *arXiv preprint arXiv:1503.02531*.
- Hsu, G.-S., Huang, W.-F., and Yap, M. H. (2019). Edge-embedded multi-dropout framework for real-time face alignment. *IEEE Access*, 8:6032–6044.
- Huang, J., Shao, X., and Wechsler, H. (1998). Face pose discrimination using support vector machines (svm). In *Proceedings. fourteenth international conference on pattern recognition (Cat. No. 98EX170)*, volume 1, pages 154–156. IEEE.
- Jones, M. and Viola, P. (2003). Fast multi-view face detection. *Mitsubishi Electric Research Lab TR-20003-96*, 3(14):2.
- Kazemi, V. and Sullivan, J. (2014). One millisecond face alignment with an ensemble of regression trees. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1867–1874.
- Kumar, A., Alavi, A., and Chellappa, R. (2017). Kepler: Keypoint and pose estimation of unconstrained faces by learning efficient h-cnn regressors. In *2017 12th IEEE international conference on automatic face & gesture recognition (fg 2017)*, pages 258–265. IEEE.
- Lathuilière, S., Juge, R., Mesejo, P., Munoz-Salinas, R., and Horaud, R. (2017). Deep mixture of linear inverse regressions applied to head-pose estimation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 4817–4825.
- Martin, M., Van De Camp, F., and Stiefelhagen, R. (2014). Real time head model creation and head pose estimation on consumer depth cameras. In *2014 2nd International Conference on 3D Vision*, volume 1, pages 641–648. IEEE.
- Meyer, G. P., Gupta, S., Frosio, I., Reddy, D., and Kautz, J. (2015). Robust model-based 3d head pose estimation. In *Proceedings of the IEEE international conference on computer vision*, pages 3649–3657.
- Mukherjee, S. S. and Robertson, N. M. (2015). Deep head pose: Gaze-direction estimation in multimodal video. *IEEE Transactions on Multimedia*, 17(11):2094–2107.
- Murphy-Chutorian, E., Doshi, A., and Trivedi, M. M. (2007). Head pose estimation for driver assistance systems: A robust algorithm and experimental evaluation. In *2007 IEEE intelligent transportation systems conference*, pages 709–714. IEEE.
- Murphy-Chutorian, E. and Trivedi, M. M. (2008). Head pose estimation in computer vision: A survey. *IEEE transactions on pattern analysis and machine intelligence*, 31(4):607–626.
- Ng, J. and Gong, S. (2002). Composite support vector machines for detection of faces across views and pose estimation. *Image and Vision Computing*, 20(5-6):359–368.
- Niyogi, S. and Freeman, W. T. (1996). Example-based head tracking. In *Proceedings of the second international conference on automatic face and gesture recognition*, pages 374–378. IEEE.
- Ranjan, R., Patel, V. M., and Chellappa, R. (2017a). Hyperface: A deep multi-task learning framework for face detection, landmark localization, pose estimation, and gender recognition. *IEEE transactions on pattern analysis and machine intelligence*, 41(1):121–135.
- Ranjan, R., Sankaranarayanan, S., Castillo, C. D., and Chellappa, R. (2017b). An all-in-one convolutional neural network for face analysis. In *2017 12th IEEE International Conference on Automatic Face & Gesture Recognition (FG 2017)*, pages 17–24. IEEE.
- Ruiz, N., Chong, E., and Rehg, J. M. (2018). Fine-grained head pose estimation without keypoints. In *Proceedings of the IEEE conference on computer vision and pattern recognition workshops*, pages 2074–2083.
- Schwarz, A., Haurilet, M., Martinez, M., and Stiefelhagen, R. (2017). Driveahead-a large-scale driver head pose

- dataset. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops*, pages 1–10.
- Seemann, E., Nickel, K., and Stiefelhagen, R. (2004). Head pose estimation using stereo vision for human-robot interaction. In *Sixth IEEE International Conference on Automatic Face and Gesture Recognition, 2004. Proceedings.*, pages 626–631. IEEE.
- Sherrah, J., Gong, S., and Ong, E.-J. (1999). Understanding pose discrimination in similarity space. In *BMVC*, pages 1–10. Citeseer.
- Sherrah, J., Gong, S., and Ong, E.-J. (2001). Face distributions in similarity space under varying head pose. *Image and Vision Computing*, 19(12):807–819.
- Srinivas, A., Lin, T.-Y., Parmar, N., Shlens, J., Abbeel, P., and Vaswani, A. (2021). Bottleneck transformers for visual recognition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 16519–16529.
- Sun, Y., Wang, X., and Tang, X. (2013). Deep convolutional network cascade for facial point detection. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3476–3483.
- Wang, Y., Liang, W., Shen, J., Jia, Y., and Yu, L.-F. (2019). A deep coarse-to-fine network for head pose estimation from synthetic data. *Pattern Recognition*, 94:196–206.
- Xie, S., Girshick, R., Dollár, P., Tu, Z., and He, K. (2017). Aggregated residual transformations for deep neural networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1492–1500.
- Xin, M., Mo, S., and Lin, Y. (2021). Eva-gcn: Head pose estimation based on graph convolutional networks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1462–1471.
- Xiong, X. and De la Torre, F. (2015). Global supervised descent method. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2664–2673.
- Yan, S., Xiong, Y., and Lin, D. (2018). Spatial temporal graph convolutional networks for skeleton-based action recognition. In *Thirty-second AAAI conference on artificial intelligence*.
- Yang, T.-Y., Chen, Y.-T., Lin, Y.-Y., and Chuang, Y.-Y. (2019). Fsa-net: Learning fine-grained structure aggregation for head pose estimation from a single image. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1087–1096.
- Yang, T.-Y., Huang, Y.-H., Lin, Y.-Y., Hsiu, P.-C., and Chuang, Y.-Y. (2018). Ssr-net: A compact soft stage-wise regression network for age estimation. In *IJCAI*, volume 5, page 7.
- Zhang, H., Wu, C., Zhang, Z., Zhu, Y., Lin, H., Zhang, Z., Sun, Y., He, T., Mueller, J., Manmatha, R., et al. (2020). Resnet: Split-attention networks. *arXiv preprint arXiv:2004.08955*.
- Zhang, Z., Hu, Y., Liu, M., and Huang, T. (2006). Head pose estimation in seminar room using multi view face detectors. In *International evaluation workshop on classification of events, activities and relationships*, pages 299–304. Springer.
- Zhou, Y. and Gregson, J. (2020). Whenet: Real-time fine-grained estimation for wide range head pose. *arXiv preprint arXiv:2005.10353*.
- Zhu, X., Lei, Z., Liu, X., Shi, H., and Li, S. Z. (2016). Face alignment across large poses: A 3d solution. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 146–155.
- Zhu, X. and Ramanan, D. (2012). Face detection, pose estimation, and landmark localization in the wild. In *2012 IEEE conference on computer vision and pattern recognition*, pages 2879–2886. IEEE.