International Journal of Research in Engineering, Science and Management
Volume 5, Issue 7, July 2022
https://www.ijresm.com | ISSN (Online): 2581-5792

Age and Gender Prediction from Facial Features Using Deep Learning

Ankita Mondal[1*], Sonam Verma[2], Sankalp Sanjay[3], Sumithra Devi[4]

[1,2,3]B.E. Student, Department of Information Science and Engineering, Dayanand Sagar Academy of Technology and Management, Bangalore, India

[4]Head of the Department, Department of Information Science and Engineering, Dayanand Sagar Academy of Technology and Management, Bangalore, India

***Abstract*: Automatic age and gender classification has been significant to a rising variety of applications since the rise of social platforms and social media. The performance of existing algorithms on real-world images, on the other hand, is pathetically inadequate, especially when compared to the massive leaps in performance recently reported for the related task of facial recognition. In this paper, we show that learning representations with deep convolutional neural networks (CNN) can result in considerable improvements in performance on certain tasks [1]. We propose a simple convolutional net design that can be used even when learning data is scarce to do it.**

***Keywords*: Cascade classifier for facial features extraction, Faster RCNN architecture, Machine Learning (ML), Convolutional Neural Network (CNN), K-Nearest Neighbor (KNN).**

1. Introduction

Age and gender are important factors in social relationships. In many languages, men and women have separate salutations and grammar norms. Different salutations and grammatical rules are reserved for men and women in distinct languages, and different vocabularies are frequently employed when addressing elders compared to young people. Despite the importance of these characteristics in our daily lives, the ability to estimate them consistently and reliably from face photos is still far from meeting the requirements of commercial applications. This is especially puzzling in light of recent claims of superhuman powers in the related job of facial recognition [2]. We aim to bridge the gap between automatic face recognition capabilities and age and gender estimation approaches in this paper. To do this, we will follow the successful lead of current facial recognition systems: Face recognition approaches reported in recent years have shown that deep convolutional neural networks (CNN) may make significant progress. We show that a basic network architecture can achieve similar results [3].

2. Related Work

Age and Gender Classification:

Calculating ratios between multiple measurements of face features was one of the first strategies for estimating age. After facial features (e.g., eyes, nose, mouth, chin, etc.) have been identified and their sizes and distances measured, ratios between them are calculated and used to categorize the face into several categories according to handcrafted rules [4]. The distribution of facial patches was represented using Gaussian Mixture Models (GMM). GMM were employed again to describe the distribution of local face data, but instead of pixel patches, robust descriptors were used [5]. Finally, super-vectors were utilized in to depict face patch distributions instead of GMM or Hidden-Markov Model.

3. Proposed System

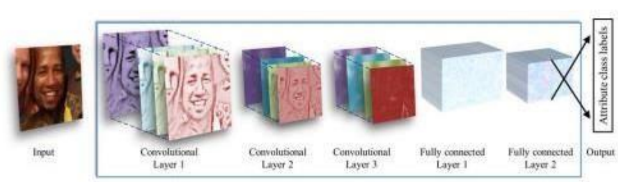

Fig. 1. Architecture of proposed system

The FERET benchmark was utilized in most of the methods outlined above, both to create the suggested systems and to evaluate their performance. FERET photos were captured in a controlled environment, making them far less problematic than face images captured in the wild. Furthermore, the results of this benchmark show that it is saturated and not difficult for modern methods to solve. As a result, estimating the real relative value of various strategies is difficult. As a result, it was tested with the widely known LFW (Labeled Faces in the Wild) benchmark, which is generally used for face recognition. A combination of LBP features and an AdaBoost classifier is used in the method [6].

4. Result Analysis

Tables 1 and 2 show our gender and age classification results, respectively. We measure and compare the accuracy of age classification when the algorithm yields the correct age group classification and when the algorithm is off by one neighboring age-group (i.e., the subject belongs to the group that is immediately older or younger than the predicted group) [7]. This follows in the footsteps of others who have done so before, and reflects the inherent unpredictability of the

*Corresponding author: mon.ankita77@gmail.com

endeavor - facial features generally vary very little between the oldest faces in one age class and the youngest faces in the next. Both tables compare results to those obtained using the procedures outlined in [8].

	0-2	4-6	8-13	15-20	25-32	38-43	48-53	60-	Total
Male	745	928	934	734	2308	1294	392	442	8192
Female	682	1234	1360	919	2589	1056	433	427	9411
Both	1427	2162	2294	1653	4897	2350	825	869	19487

Table 1. **The AdienceFaces benchmark.** Breakdown of the AdienceFaces benchmark into the different Age and Gender classes.

Method	Accuracy
Best from [10]	77.8 ± 1.3
Best from [23]	79.3 ± 0.0
Proposed using single crop	85.9 ± 1.4
Proposed using over-sample	$\mathbf{86.8 \pm 1.4}$

Table 2. Gender estimation results on the Adience benchmark. Listed are the mean accuracy±standard error over all age categories. Best results are marked in bold.

On both tasks, the suggested technique clearly outperforms the reported state- of-the-art with significant gaps. The contribution of the over-sampling strategy, which delivers a performance advantage over the original network, is also visible. This suggests that greater alignment (for example, frontalization) may provide a performance gain. Figures 1 and 2 show a few examples of gender and age misclassifications, respectively. These results reveal that the extremely challenging viewing conditions of several of the audience benchmark photographs are the root of many of our system's faults. The most obvious faults are those brought on by blur, poor resolution, and occlusions (especially with thick makeup. In pictures of newborn or young children, where obvious gender features are not yet visible, gender assessment errors are also frequent [9].

5. Conclusion

Although many prior techniques addressed the challenges of categorising people by age and gender, much of this work has focused on the few photographs taken in lab settings. The extremely challenging viewing conditions of several of the Audience benchmark photographs are the root of many of our system's faults.

The most obvious faults are those brought on by blur, poor resolution, and occlusions (especially with thick makeup). In pictures of newborns or very young children, where obvious gender features are not yet visible, gender assessment errors are also frequent.

On the other hand, are not only more difficult, but also more plentiful. The easy availability of huge image collections provides modern machine learning based systems with effectively endless training data, though this data is not always suitably labeled for supervised learning.

Using Internet data and the associated issue of facial recognition as an example, we examine how well deep CNN performs on these tasks. Due to the dearth of labelled data, we avoid over fitting by employing a lightweight deep-learning architecture to display our findings. Our network is "thin" in compared to certain modern network architectures, reducing the amount of parameters and the chance of overfitting.

To increase the amount of the training data, we purposefully include cropped copies of the images in our training set. When tested on the Audience benchmark of unfiltered pictures, the resulting system performed noticeably better than the state of the art.

Our findings lead to two crucial conclusions. First, even with the significantly lower size of today's unconstrained image sets labelled for age and gender, CNN may be used to improve age and gender classification results. Second, because our model is so simple, more complex systems with additional training data may potentially be capable of much bettering the findings given here [10].

References

[1] T. Ahonen, A. Hadid, and M. Pitkanen. Face recognition using local binary patterns for face description. 2 Trans. Pattern Anal. Mach. Intell., 28(12):2:2037– 2041, 2006.2 Trans. Pattern Anal. Mach. Intell., 28(12): 2037–2041, 2006.

[2] S. Baluja and H. A. Rowley are the authors of this article. Increasing the accuracy of sex identification. International Journal of Computer Vision, 71(1), 111–119, 2007.

[3] A. Bar-Hillel, T. Hertz, N. Shental, and D. Weinshall. Using equivalence relations to learn distance functions. In volume 3 of the International Conference on Machine Learning, pages 11–18, 2003.

[4] W.-L. Chao, J.-Z. Liu, and J.-J. Ding. Label-sensitive learning and age-oriented regression are used to estimate facial age. Pattern Recognition, vol. 46, no. 3, pp. 628–641, 2013.

[5] K. Chatfield, K. Simonyan, A. Vedaldi, and A. Zisserman. Return of the devil in the details: Delving deep into convolutional nets. arXiv preprint arXiv:1405.3531, 2014.

[6] J. Chen, S. Shan, C. He, G. Zhao, M. Pietikainen, X. Chen, and W. Gao. Wld: A local image descriptor with a lot of power. 32(9):1705–1720, Trans. Pattern Anal. Mach. Intell., 2010.

[7] T. F. Cootes, G. J. Edwards, and C. J. Taylor. Active appearance models. In European Conf. Comput. Vision, pages 484–498. Springer, 1998. 2

[8] C. Cortes and V. Vapnik. Support-vector networks. Machine learning, 20(3):273–297, 1995.

[9] E. Eidinger, R. Enbar, and T. Hassner. Age and gender estimation of unfiltered faces. Trans. On Inform. Forensics and Security, 9(12), 2014.

[10] Y. Fu, G. Guo, and T. S. Huang. Age synthesis and estimation via faces: A survey. Trans. Pattern Anal. Mach. Intell., 32(11):1955–1976, 2010.