

Age and Gender Classification using Convolutional Neural Networks

Adabala Bhaskar Sai, Durga Prasad

Pg Scholar, Department Of Mca, Dnr College, Bhimavaram, Andhra Pradesh.

K. Rambabu

(Assistant Professor), Master Of Computer Applications, Dnr college, Bhimavaram, Andhra Pradesh.

Abstract: *Age and gender play fundamental roles in social interactions. Languages reserve different salutations and grammar rules for men or women, and very often different vocabularies are used when addressing elders compared to young people. In this paper we show that by learning representations through the use of deep-convolutional neural networks (CNN), a significant increase in performance can be obtained on these tasks. To this end, we propose a simple convolutional net architecture that can be used even when the amount of learning data is limited. We evaluate our method on the recent Adience benchmark for age and gender estimation and show it to dramatically outperform current state-of-the-art methods.*

Keywords: *CNN, age, gender.*

I. INTRODUCTION

Age and gender play fundamental roles in social interactions. Languages reserve different salutations and grammar rules for men or women, and very often different vocabularies are used when addressing elders compared to young people. Despite the basic roles these attributes play in our day-to-day lives, the ability to automatically estimate them accurately and reliably from face images is still far from meeting the needs of commercial applications.

This is particularly perplexing when considering recent claims to super-human capabilities in the related task of face recognition. Past approaches to estimating or classifying these attributes from face images have relied on differences in facial feature dimensions or “tailored” face descriptors. Most have employed classification schemes designed particularly for age or gender estimation tasks, including and others. Few of these past methods were designed to handle the many challenges of unconstrained imaging conditions. Moreover, the machine learning methods employed by these systems did not fully exploit the massive numbers of image examples and data available through the

Internet in order to improve classification capabilities.

In this paper we attempt to close the gap between automatic face recognition capabilities and those of age and gender estimation methods. To this end, we follow the successful example laid down by recent face recognition systems: Face recognition techniques described in the last few years have shown that tremendous progress can be made by the use of deep convolutional neural networks (CNN). We demonstrate similar gains with a simple network architecture, designed by considering the rather limited availability of accurate age and gender labels in existing face data sets. We test our network on the newly released Adience benchmark for age and gender classification of unfiltered face images. We show that despite the very challenging nature of the images in the Adience set and the simplicity of our network design, our method outperforms existing state of the art by substantial margins. Although these results provide a remarkable baseline for deep-learning-based approaches, they leave room for improvements by more elaborate system designs, suggesting that the problem of accurately estimating age and gender in the unconstrained settings, as reflected by the Adience images, remains unsolved. In order to provide a foothold for the development of more effective future methods, we make our trained models and classification system publicly available.

1.1 Objective of the Project

Age and gender play fundamental roles in social interactions. Languages reserve different salutations and grammar rules for men or women, and very often different vocabularies are used when addressing elders compared to young people. In this paper we show that by learning representations through the use of deep-convolutional neural networks (CNN), a significant increase in

performance can be obtained on these tasks. To this end, we propose a simple convolutional net architecture that can be used even when the amount of learning data is limited. We evaluate our method on the recent Adience benchmark for age and gender estimation and show it to dramatically outperform current state-of-the-art methods.

II. LITEARTURE SURVEY

Face description with local binary patterns: Application to face recognition.

This paper presents a novel and efficient facial image representation based on local binary pattern (LBP) texture features. The face image is divided into several regions from which the LBP feature distributions are extracted and concatenated into an enhanced feature vector to be used as a face descriptor. The performance of the proposed method is assessed in the face recognition problem under different challenges. Other applications and several extensions are also discussed.

Boosting sex identification performance

This paper presents a method based on AdaBoost to identify the sex of a person from a low resolution grayscale picture of their face. The method described here is implemented in a system that will process well over 109 images. The goal of this work is to create an efficient system that is both simple to implement and maintain; the methods described here are extremely fast and have straightforward implementations. We achieve 80% accuracy in sex identification with less than 10 pixel comparisons and 90% accuracy with less than 50 pixel comparisons. The best classifiers published to date use Support Vector Machines; we match their accuracies with as few as 500 comparison operations on a 20×20 pixel image. The AdaBoost based classifiers presented here achieve over 93% accuracy; these match or surpass the accuracies of the SVM-based classifiers, and yield performance that is 50 times faster.

Learning distance functions using equivalence relations

We address the problem of learning distance metrics using side-information in the form of groups of "similar" points. We propose to use the

RCA algorithm, which is a simple and efficient algorithm for learning a full ranked Mahalanobis metric (Shental *et al.*, 2002). We first show that RCA obtains the solution to an interesting optimization problem, founded on an information theoretic basis. If the Mahalanobis matrix is allowed to be singular, we show that Fisher's linear discriminant followed by RCA is the optimal dimensionality reduction algorithm under the same criterion. We then show how this optimization problem is related to the criterion optimized by another recent algorithm for metric learning (Xing *et al.*, 2002), which uses the same kind of side information. We empirically demonstrate that learning a distance metric using the RCA algorithm significantly improves clustering performance, similarly to the alternative algorithm. Since the RCA algorithm is much more efficient and cost effective than the alternative, as it only uses closed form expressions of the data, it seems like a preferable choice for the learning of full rank Mahalanobis distances.

Facial age estimation based on label-sensitive learning and age-oriented regression.

In this paper, a new age estimation framework considering the intrinsic properties of human ages is proposed, which improves the dimensionality reduction techniques to learn the connections between facial features and aging labels. To enhance the performance of dimensionality reduction, a distance metric adjustment step is introduced in advance to achieve a suitable metric in the feature space. In addition, to further exploit the ordinal relationship of human ages, the "label-sensitive" concept is proposed, which regards the label similarity during the learning phase of distance metric and dimensionality reduction. Finally, an age-specific local regression algorithm is proposed to capture the complicated aging process for age determination. From the simulation results, the proposed framework achieves the lowest mean absolute error against the existing methods.

Human age estimation with regression on discriminative aging manifold.

Recently, extensive studies on human faces in the human-computer interaction (HCI) field reveal

significant potentials for designing automatic age estimation systems via face image analysis. The success of such research may bring in many innovative HCI tools used for the applications of human-centered multimedia communication. Due to the temporal property of age progression, face images with aging features may display some sequential patterns with low-dimensional distributions. In this paper, we demonstrate that such aging patterns can be effectively extracted from a discriminant subspace learning algorithm and visualized as distinct manifold structures. Through the manifold method of analysis on face images, the dimensionality redundancy of the original image space can be significantly reduced with subspace learning. A multiple linear regression procedure, especially with a quadratic model function, can be facilitated by the low dimensionality to represent the manifold space embodying the discriminative property. Such a processing has been evaluated by extensive simulations and compared with the state-of-the-art methods. Experimental results on a large size aging database demonstrate the effectiveness and robustness of our proposed framework.

Face Age Classification on Consumer Images with Gabor Feature and Fuzzy LDA Method

As we all know, face age estimation task is not only challenging for computer, but even hard for human in some cases, however, coarse age classification such as classifying human face as baby, child, adult or elder people is much easier for human. In this paper, we try to dig out the potential age classification power of computer on faces from consumer images which are taken under various conditions. Gabor feature is extracted and used in LDA classifiers. In order to solve the intrinsic age ambiguity problem, a fuzzy version LDA is introduced through defining age membership functions. Systematic comparative experiment results show that the proposed method with Gabor feature and fuzzy LDA can achieve better age classification precision in consumer images.

III. PROPOSED METHOD

In this paper, we proposed a age and gender predicted framework by using deep-convolutional neural networks (CNN). Two important conclusions can be made from our results. First,

CNN can be used to provide improved age and gender classification results, even considering the much smaller size of contemporary unconstrained image sets labeled for age and gender. Second, the simplicity of our model implies that more elaborate systems using more training data may well be capable of substantially improving results beyond those reported here.

Advantages:

1. Accuracy is high.
2. Simple architecture.

IV. RESULTS




Fig. Age and gender detected



Fig. Age and gender detected

V. CONCLUSION

Though many previous methods have addressed the problems of age and gender classification, until recently, much of this work has focused on constrained images taken in lab settings. Such settings do not adequately reflect appearance variations common to the real-world images in

social websites and online repositories. Internet images, however, are not simply more challenging: they are also abundant. 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. Taking example from the related problem of face recognition we explore how well deep CNN perform on these tasks using Internet data. We provide results with a lean deep-learning architecture designed to avoid overfitting due to the limitation of limited labeled data. Our network is “shallow” compared to some of the recent network architectures, thereby reducing the number of its parameters and the chance for overfitting. We further inflate the size of the training data by artificially adding cropped versions of the images in our training set. The resulting system was tested on the Adience benchmark of unfiltered images and shown to significantly outperform recent state of the art. Two important conclusions can be made from our results. First, CNN can be used to provide improved age and gender classification results, even considering the much smaller size of contemporary unconstrained image sets labeled for age and gender. Second, the simplicity of our model implies that more elaborate systems using more training data may well be capable of substantially improving results beyond those reported here.

REFERENCES

- [1] T. Ahonen, A. Hadid, and M. Pietikainen. Face description with local binary patterns: Application to face recognition. *Trans. Pattern Anal. Mach. Intell.*, 28(12):2037–2041, 2006. 2
- [2] S. Baluja and H. A. Rowley. Boosting sex identification performance. *Int. J. Comput. Vision*, 71(1):111–119, 2007. 2
- [3] A. Bar-Hillel, T. Hertz, N. Shental, and D. Weinshall. Learning distance functions using equivalence relations. In *Int. Conf. Mach. Learning*, volume 3, pages 11–18, 2003. 2
- [4] W.-L. Chao, J.-Z. Liu, and J.-J. Ding. Facial age estimation based on label-sensitive learning and age-oriented regression. *Pattern Recognition*, 46(3):628–641, 2013. 1, 2
- [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. 3

- [6] J. Chen, S. Shan, C. He, G. Zhao, M. Pietikainen, X. Chen, and W. Gao. Wld: A robust local image descriptor. *Trans. Pattern Anal. Mach. Intell.*, 32(9):1705–1720, 2010. 2

- [7] S. E. Choi, Y. J. Lee, S. J. Lee, K. R. Park, and J. Kim. Age estimation using a hierarchical classifier based on global and local facial features. *Pattern Recognition*, 44(6):1262–1281, 2011. 2

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

- [9] C. Cortes and V. Vapnik. Support-vector networks. *Machine learning*, 20(3):273–297, 1995. 2

- [10] E. Eidinger, R. Enbar, and T. Hassner. Age and gender estimation of unfiltered faces. *Trans. on Inform. Forensics and Security*, 9(12), 2014. 1, 2, 5, 6