Identifying Gender from Unaligned Facial Images by Set Classification

Wen-Sheng Chu[1], Chun-Rong Huang[1] and Chu-Song Chen[1,2,3]
{wschu,nckuos,song}@iis.sinica.edu.tw

[1]Institute of Information Science, Academia Sinica, Taipei, Taiwan
[2]Research Center for Information Technology Innovation, Academia Sinica, Taipei, Taiwan
[3]Graduate Institute of Networking and Multimedia, National Taiwan University, Taipei, Taiwan

August 25, 2010

Introduction

- Gender classification: determine the sex of a person from a face image.

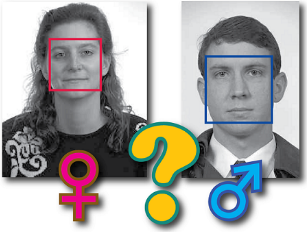

- A useful preprocessing step for face applications
  - Pre-separate face images into two partitions
  - Examples: face recognition, age estimation, facial expression recognition. . .
- Other interesting applications
  - psychology
  - security industry
  - human-computer interaction
  - demographic data collection

Outline

1. Face Alignment v.s. Gender Classification
2. The Proposed Approach
3. Experimental Results
4. Conclusion

Outline

1. Face Alignment v.s. Gender Classification
2. The Proposed Approach
3. Experimental Results
4. Conclusion

Face Alignment v.s. Gender Classification

- Gender classification pipeline

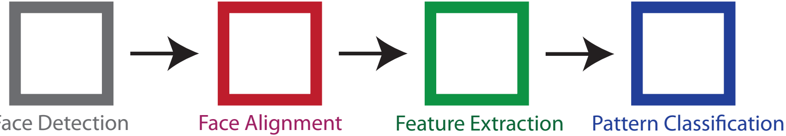

- Face detection could not specify a robust face location.
  - In practice, false recognition may occur frequently due to variations of human poses or difficult alignments.
- Face alignment methods could reduce the variability during the training phase for gender classification[1].
  - But, facial landmark detectors must be accurate and manual efforts are required to label specific facial features.

[1] E. Mäkinen and R. Raisamo. Evaluation of gender classification methods with automatically detected and aligned faces. *TPAMI*, 30(3):541–547, 2008a. doi: 10.1109/TPAMI.2007.70800

Related Work

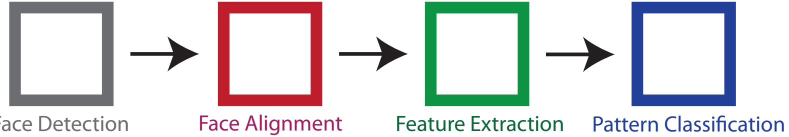

- Feature Extraction: Gabor wavelets, haar-like features, local binary patterns, active appearance model, raw pixel intensity...
- Pattern Classification: neural networks, support vector machines, boosting, linear discriminant analysis, genetic algorithms...

Lyons et al. (2000), AFGR
Shakhnarovich et al. (2002), PAMI
Moghaddam and Yang (2002), AFGR
Sun et al. (2002), WACV
Wu et al. (2003), AVBPA
Jain and Huang (2004), AFGR
Costen et al. (2004), AFGR
Lian and Lu (2006), ISNN
Baluja and Rowley (2007), IJCV
Mäkinen and Raisamo (2008a), PAMI
Mäkinen and Raisamo (2008b), PR Letters

Potential Problems of Face Alignment Methods

- Typical face alignment methods:
  - Geometric shape alignment: affine warping
  - Statistical model: ASM/AAM (Cootes et al., 2001)

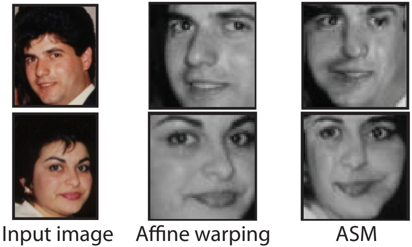

- Potential problems
  1. Affine warping changes the geometric ratio of faces.
  2. ASM/AAM assumes the transformation are used for frontal faces.
  3. ASM/AAM needs to rebuild the model for every new dataset.
  4. Facial landmark detectors must be accurate.
  5. Rough alignments downgrade the performance.

To what extent can face alignment really help gender classification?

Learning from Sets

- How about using an image set instead of just one input image?

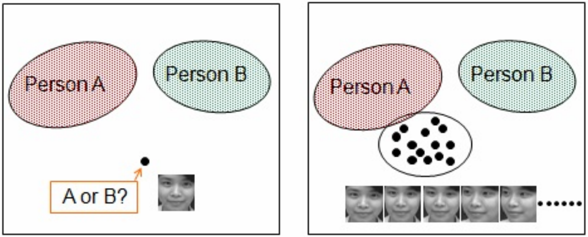

- An image set represents more variation in an object's appearance, such as changes of camera poses and lighting conditions.
- More robust recognition should be achieved by set information rather than just a single input.

Outline

1. Face Alignment v.s. Gender Classification

2. The Proposed Approach

3. Experimental Results

4. Conclusion

Our Idea

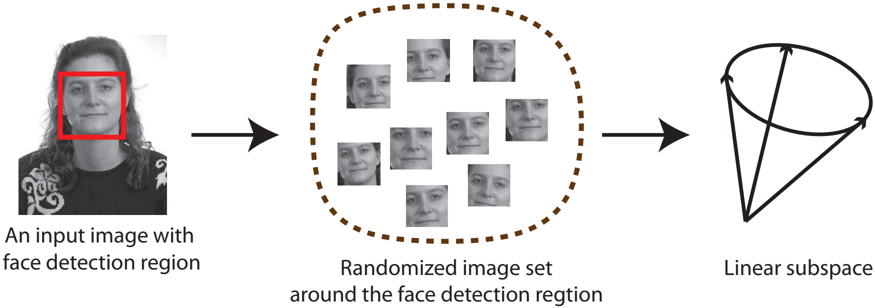
An input image with
face detection region
Randomized image set
around the face detection regtion
Linear subspace

The Proposed Approach

- Goal: to avoid alignment processes and to discover the gender correlation from unaligned facial image sets.
- "How to align well?"⇒"How to avoid alignment?"
- We collect a set of facial patches according to random scaling and rotation in a neighborhood of the face detection region.

- The concept of simulating invariances of interest through synthetic copies is similar to "kernel jittering[2]".

[2] D. Decoste and B. Schölkopf. Training invariant support vector machines. *Machine Learning*, 46(1):161–190, 2002

Representation for an Image Set

- We collected 75 facial images as the image set for each subject.
- Rather than directly using the scattered distribution, we represent each image set $\mathbf{X}_i$ as a linear subspace $\mathcal{L}_i$.
- Given image sets $\{\mathbf{X}_i, \ldots, \mathbf{X}_N\}$ for N subjects, an unitary basis matrix $\mathbf{P}_i$ for $\mathcal{L}_i$ can be obtained by the singular value decomposition (SVD)

$$\mathbf{X}_i \mathbf{X}_i^T \simeq \mathbf{P}_i \Lambda \mathbf{P}_i^T.$$

- A representation of linear subspace derives the following benefits:
  1. capture the joint information (common parts) in the set of facial patches
  2. deal with the problem of outliers and missing pixels (occlusion)
  3. retain the property of approximating the original dataset well

Similarity between Two Subspaces

- Idea: measure the similarity as the closeness of the closest vectors from two subspaces.
- This is known as the canonical correlations, which can be uniquely defined as the cosines of principal angles between basis vectors.
- Principal angles can be obtained from the SVD solutions[3] of two d-dimensional unitary orthogonal bases $\mathbf{P}_i$ and $\mathbf{P}_j$ for linear subspaces $\mathcal{L}_i$ and $\mathcal{L}_j$:

$$\mathbf{P}_i^T \mathbf{P}_j = \mathbf{U}_{ij} \Lambda \mathbf{U}_{ji}^T \;\; s.t. \;\; \Lambda = diag(\sigma_1, \ldots, \sigma_d),$$

where the singular values $0 \leq \sigma_i \leq 1$ are defined as canonical correlations.

[3] Å. Björck and G. H. Golub. Numerical methods for computing angles between linear subspaces. *Mathematics of Computation*, 27:579–594, 1973

Discriminative Learning for Gender Classification

- Classical linear discriminant analysis seeks for a transformation such that the separation between genders is maximized and that within genders is minimized.

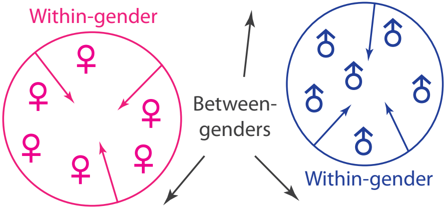

- Scatter measures of within- and between-class can be related to pairwise correlation by the scatter matrices:

$$\mathbf{B} = \sum_i \mathbf{x}_i^T \mathbf{x}_i^B, \quad \mathbf{W} = \sum_i \mathbf{x}_i^T \mathbf{x}_i^W,$$

where $\mathbf{x}_i^B$, $\mathbf{x}_i^W$ indicate the closet between-gender and within-gender image vectors of a given $\mathbf{x}_i$.

Gender Discriminant Transformation

- The aim is to find a transformation $\mathbf{T}$ that transform all subspaces to another space containing maximal correlation within the same gender and minimal correlation between different genders.
- To find $\mathbf{T}$, we employ the discriminant analysis of canonical correlations (DCC)[4] by solving

$$\mathbf{T} = \arg\max_{\mathbf{T}} \frac{\mathrm{tr}(\mathbf{T}^T \mathbf{S}_b \mathbf{T})}{\mathrm{tr}(\mathbf{T}^T \mathbf{S}_w \mathbf{T})},$$

  where $\mathbf{S}_b$ and $\mathbf{S}_w$ are the between-gender and within-gender scatter matrices.
- The transformation matrix $\mathbf{T}$ is solved as the eigen-decomposition of $\mathbf{S}_w^{-1}\mathbf{S}_b$.

[4] T.-K. Kim, J. Kittler, and R. Cipolla. Discriminative learning and recognition of image set classes using canonical correlations. *TPAMI*, 29(6):1005–1018, 2007. doi: 10.1109/TPAMI.2007.1037

Outline

1. Face Alignment v.s. Gender Classification
2. The Proposed Approach
3. Experimental Results
4. Conclusion

Datasets

- We experimented the proposed approach on two databases.

(a) FERET[5]

(b) MORPH[6]

- MORPH is a more challenging dataset with more images than FERET in terms of ages, facial poses and expressions.

[5] P. J. Phillips, H. Wechsler, J. Huang, and P. J. Rauss. The FERET database and evaluation procedure for face-recognition algorithms. *Image and Vision Computing*, 16(5):295–306, 1998

[6] M. Minear and D. C. Park. A lifespan database of adult facial stimuli. *Behavior Research Methods Instruments and Computers*, 36(4):630–633, 2004

Experiment Setups

- The number of training female images and male images are equivalent.
- To compare our results to those reported in Mäkinen and Raisamo (2008b), we used the same setting for the FERET database as follows.
  - There are 450 female and 450 male images.
  - Each subject contains only one image.
- For the MORPH database, we collected 8,033 female and 47,810 male images.
  - Totally 418 subjects
  - Age ranging from 18-69
  - Races including Caucasian, African-American and Asian

Experimental Results

- Comparison between two set-based methods, DCC[7] and MSM[8], and sample-based methods.

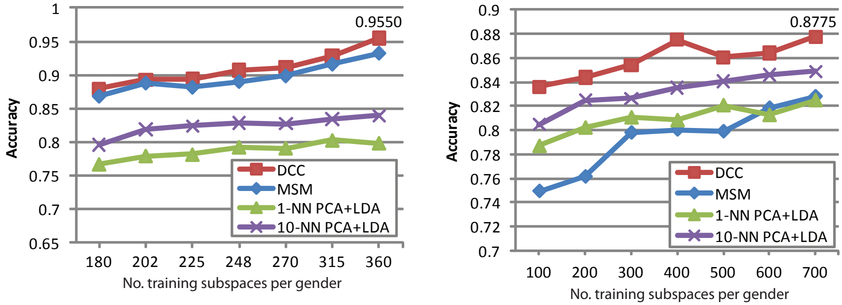

(a) FERET (b) MORPH

[7] T.-K. Kim, J. Kittler, and R. Cipolla. Discriminative learning and recognition of image set classes using canonical correlations. *TPAMI*, 29(6):1005–1018, 2007. doi: 10.1109/TPAMI.2007.1037

[8] O. Yamaguchi, K. Fukui, and K. I. Maeda. Face recognition using temporal image sequence. In *IEEE International Conference on Automatic Face and Gesture Recognition*, pages 318–323, 1998. doi: 10.1109/AFGR.1998.670968

Comparison with the State-of-art Methods

- We compared our performance with the state-of-the-art methods discussed in (Mäkinen and Raisamo, 2008b)[9].

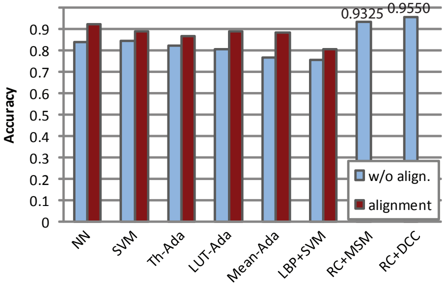

- The performance of the proposed is shown to outperform the methods without proper face alignments.

---

[9] E. Mäkinen and R. Raisamo. An experimental comparison of gender classification methods. *Pattern Recogn. Lett.*, 29(10):1544–1556, 2008b. ISSN 0167-8655. doi: http://dx.doi.org/10.1016/j.patrec.2008.03.016

How about Other Applications?

- How about other appearance variations, except for poses?
  - illumination
  - facial expression
  - ...
- How about other input patterns, except for face images?
  - multi-view object images
  - feature vectors
  - ...
- How about other tasks, except for gender classification?
  - face recognition
  - age estimation
  - facial expression recognition
  - object recognition
  - ...

Outline

1. Face Alignment v.s. Gender Classification
2. The Proposed Approach
3. Experimental Results
4. Conclusion

Conclusion

- We have presented an approach for identifying genders from unaligned facial images.
- A set of unaligned facial patches instead of a single image is taken as input.
- Each image set is represented as a linear subspace.
- Discriminative gender information is learned across the linear subspaces using set classification methods.
- The results was shown to outperform the state-of-art methods for gender classification.

References

S. Baluja and H. Rowley. Boosting sex identification performance. *International Journal of Computer Vision*, 71(1):111–119, 2007.

Å. Björck and G. H. Golub. Numerical methods for computing angles between linear subspaces. *Mathematics of Computation*, 27:579–594, 1973.

T. F. Cootes, G. J. Edwards, and C. J. Taylor. Active appearance models. *TPAMI*, 23(6):681, 2001.

N. Costen, M. Brown, and S. Akamatsu. Sparse models for gender classification. In *Proceedings of the 6th International Conference on Automatic Face and Gesture Recognition*, 2004.

D. Decoste and B. Schölkopf. Training invariant support vector machines. *Machine Learning*, 46(1):161–190, 2002.

A. Jain and J. Huang. Integrating independent components and linear discriminant analysis for gender classification. In *Sixth IEEE International Conference on Automatic Face and Gesture Recognition, 2004. Proceedings*, pages 159–163, 2004.

T.-K. Kim, J. Kittler, and R. Cipolla. Discriminative learning and recognition of image set classes using canonical correlations. *TPAMI*, 29(6):1005–1018, 2007. doi: 10.1109/TPAMI.2007.1037.

H. C. Lian and B. L. Lu. Multi-view gender classification using local binary patterns and support vector machines. *Internat. Sympos. on Neural Networks*, 3972:202, 2006.

M. J. Lyons, J. Budynek, A. Plante, and S. Akamatsu. Classifying facial attributes using a 2-D Gabor waveletrepresentation and discriminant analysis. In *Automatic Face and Gesture Recognition*, pages 202–207, 2000.

E. Mäkinen and R. Raisamo. Evaluation of gender classification methods with automatically detected and aligned faces. *TPAMI*, 30(3):541–547, 2008a. doi: 10.1109/TPAMI.2007.70800.

E. Mäkinen and R. Raisamo. An experimental comparison of gender classification methods. *Pattern Recogn. Lett.*, 29(10): 1544–1556, 2008b. ISSN 0167-8655. doi: http://dx.doi.org/10.1016/j.patrec.2008.03.016.

M. Minear and D. C. Park. A lifespan database of adult facial stimuli. *Behavior Research Methods Instruments and Computers*, 36(4):630–633, 2004.

B. Moghaddam and M.-H. Yang. Learning gender with support faces. *TPAMI*, 24(5):707–711, 2002. doi: 10.1109/34.1000244.

P. J. Phillips, H. Wechsler, J. Huang, and P. J. Rauss. The FERET database and evaluation procedure for face-recognition algorithms. *Image and Vision Computing*, 16(5):295–306, 1998.

G. Shakhnarovich, P. A. Viola, and B. Moghaddam. A unified learning framework for real time face detection and classification. In *Automatic Face and Gesture Recognition*, pages 14–21, 2002. doi: 10.1109/AFGR.2002.1004124.

Z. Sun, G. Bebis, X. Yuan, and S. J. Louis. Genetic feature subset selection for gender classification: A comparison study. In *IEEE Workshop on Applications of Computer Vision*, pages 165–170, 2002.

M. Varma and B. Babu. More generality in efficient multiple kernel learning. In *Proceedings of the 26th Annual International Conference on Machine Learning*, pages 1065–1072. ACM, 2009.

B. Wu, H. Ai, and C. Huang. Lut-based adaboost for gender classification. *Internat. Conference on Audio and Video-based Biometric Person Authentication*, 2688:104–110, 2003.

O. Yamaguchi, K. Fukui, and K. I. Maeda. Face recognition using temporal image sequence. In *IEEE International Conference on Automatic Face and Gesture Recognition*, pages 318–323, 1998. doi: 10.1109/AFGR.1998.670968.

Thank you for your attendance!

Questions?