

Everything is in the Face? Represent Faces with Object Bank

Xin Liu^{1,2}, Shiguang Shan¹, Shaoxin Li^{1,2}, Alexander G. Hauptmann³

¹Key Lab of Intelligent Information Processing of Chinese Academy of Sciences (CAS), Institute of Computing Technology, CAS, Beijing, 100190, China

²University of Chinese Academy of Sciences, Beijing 100049, China

³School of Computer Science, Carnegie Mellon University, PA 15213, USA
{xin.liu, shiguang.shan, shaoxin.li}@vipl.ict.ac.cn, alex@cs.cmu.edu;

Abstract. Object Bank (OB) [1] has been recently proposed as an object-level image representation for high-level visual recognition. OB represents an image from its responses to many pre-trained object filters. While OB has been validated in general image recognition tasks, it might seem ridiculous to represent a face with OB. However, in this paper, we study this anti-intuitive potential and show how OB can well represent faces amazingly, which seems a proof of the saying that “Everything is in the face”. With OB representation, we achieve results better than many low-level features and even competitive to state-of-the-art methods on LFW dataset under unsupervised setting. We then show how we can achieve state of the art results by combining OB with some low-level feature (e.g. Gabor).

1 Introduction

To achieve accurate automatic face recognition, faces must be modeled appropriately. In the past decades, faces are widely represented by computer vision researchers with hand-crafted local features, e.g., Gabor [2], Local Binary Pattern (LBP) [3] and its high dimensional variant [4], patterns of oriented edge magnitudes (POEM) [5], Local Quantized Pattern (LQP) [6], Scale-Invariant Feature Transform (SIFT) [7], and Histogram of Oriented Gradients (HOG) [8], etc. However, designing an effective local descriptor demands considerable domain specific knowledge and a great deal of efforts.

Recently, data-driven representation learning is becoming popular and reports promising accuracy. In [9] and [10], filters are learned to maximize the discriminative power for face recognition. In [11], a multi-layer face representation framework is constructed based on large-scale random filter searching. In [12], [13] and [14], codebook learning technologies are utilized for robust face representation design. More recently, faces are represented with mid-level or high-level semantic information. For instance, the attributes and simile classifier [15] represents faces by the mid-level face attributes and so-called simile feature. And the Tom-versus-Pete classifier [16] encodes faces with high-level semantic information by the output scores of a large number of person-pair classifiers. In

the last two years, deep learning methods are applied to face recognition to learn hierarchical face representation and report state-of-the-art results [17, 18].

On the other hand, general object recognition has also been studied extensively in the past decade [19, 20] and has achieved significant progress recently [21]. Actually, in terms of representation, general object recognition roughly shares similar conceptual paradigms to those for face recognition, i.e., low-level local features [22–24], mid-level features [25], high-level semantic features [1], and their combination via hierarchical deep models [21]. Especially, many low-level features, such as SIFT and HOG, have been applied successfully to both face recognition [7, 8] and general object recognition [22, 26]. However, it is unclear whether the same mid-level features and high-level features can be general enough to boost both face recognition and general object recognition.

Intuitively, it seems ridiculous to apply mid-level attribute or high-level semantic features for general object representation to face recognition. This intuition is further strengthened by the findings in the cognitive neuroscience literature that human vision system seems to have specific vision cortices for face recognition, separated from those for general object recognition. There is a lot of evidence that the primary locus for human face perception is found on the fusiform gyrus (so called fusiform face area, abbr. as FFA) of the extra-striate visual cortex [27, 28]. In other words, perception mechanism of face recognition should be different from general object recognition. At least, the high-level semantic features for general object recognition are probably not sharable to face recognition, as there is indeed neither truck nor ship on human faces. Thus, computationally, it will seem silly to represent a face with general object filters like OB [1].

However, in this paper, we reach an anti-intuitive observation that faces can be computationally represented by high-level general “object filters”, which seems to validate the saying that “everything is in the face”. Specifically, we model faces by encoding the responses of OB filters operated on the input faces. With this OB representation, we surprisingly achieve results better than many low-level features and even competitive to state-of-the-art methods on LFW dataset [29] under unsupervised setting. We also investigate the influence of pooling methods, the scales of object filters, and importance of various objects. Finally, we show how we can achieve state-of-the-art results by combining OB with some low-level features (e.g. Gabor).

Our finding, i.e., faces can be represented by high-level object-based models, seems incredible at first glance. Nevertheless, we argue that it might be actually consistent with some cognitive observations, which supports the argument that faces are not processed specifically different from general object recognition [30]. It has been suggested that rather than being a true face module, the FFA may be responsible for performing either subordinate or expert-level categorization of generic objects, as suggested from both behavioral studies [31, 32] and neuroimaging studies [33]. However, we must admit that it is still too early to make any determinate decision on what our findings suggest.

The rest of this paper is organized as follows: section 2 briefly reviews OB and presents its application to face recognition. Then, in section 3, we evaluate its accuracy on LFW restricted face recognition evaluation under unsupervised setting. The following section then describes how performance can be further improved by combining OB with other low-level features. Finally, we conclude the work and discuss future work in the last section.

2 Object Bank Face: Face Representation with Object Bank

In this section, we firstly briefly review the OB method [1]. Then, we describe how we apply OB to represent faces for face recognition.

2.1 Object Bank

OB is a high-level image representation which explicitly encodes the appearance and spatial location information of object in the image. OB represents image by collecting its responses to a large number of pre-trained object filters.

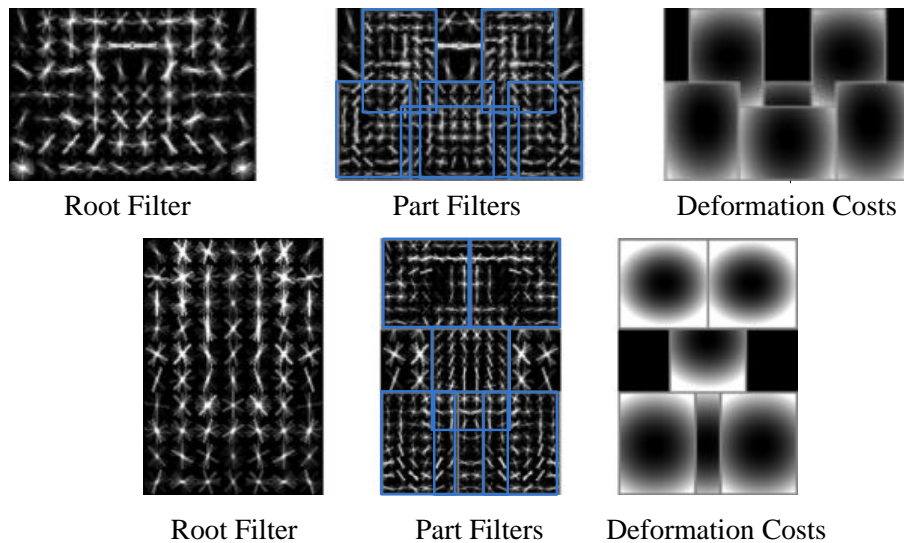


Fig. 1. An example DPM of bicycle trained on ImageNet dataset [20]. Two components are demonstrated, each component consists of a coarse root filter (1), multiple part filters (2) and a spatial model indicates the relative position of each part (3). The upper row shows the first component, which captures sideways views of bicycles. The lower row shows the second component, which captures frontal and near frontal views of bicycles.

In the OB framework, Deformable Part Model (DPM) [26] is employed as the object detector/filter. DPM is a widely used method for learning multi-view object detectors, and it is able to discriminatively learn a part-based object detector model using training samples only with the bounding box labels. This is done by treating the object parts as latent variables and learning by latent SVM method [26]. Each DPM is defined by a coarse root filter with multiple part filters. Fig. 1 shows an example DPM of bicycle.

In Fig. 2, we demonstrate the feature extraction procedure of OB, including two main steps: convolution and pooling.

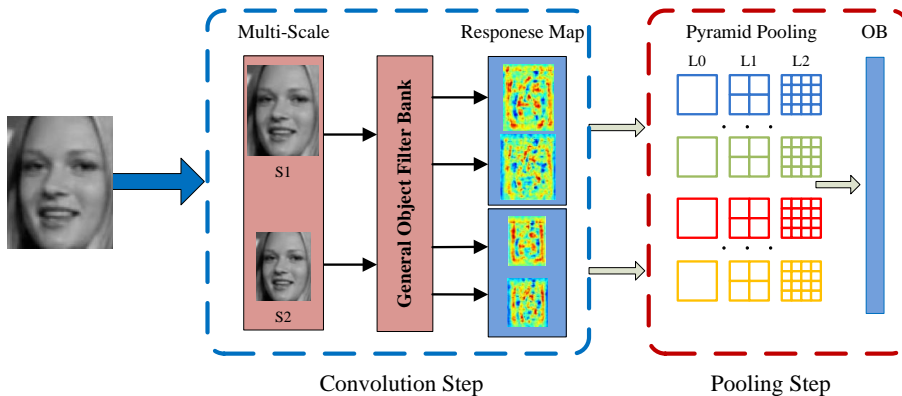


Fig. 2. Illustration of the OB. The convolution step is encircled by blue dotted box and the pooling step is encircled by dark red dotted box.

In the convolution step, pre-trained object filters are convolved with a multi-scale image pyramid of the input image. The response maps of all object filters in different scales are collected all together. Note that each object filter has two sets of filters with respect to different viewpoints, i.e., frontal and profile. So, in each scale, there are two response maps for each object filter. The final number of response map is $NumObject * NumScale * 2$. The value of each location on the response map indicates how likely an image patch contains a specific object.

In the pooling step, OB encodes both the appearance and spatial location information of objects. OB employs a 3-level pyramid pooling strategy. Each response map is divided into $1 * 1$ (L_0), $2 * 2$ (L_1) and $4 * 4$ (L_2) grids. Two kinds of pooling operators are considered in this work. The first is the max-pooling, which computes the maximum response value in each grid. The second is the mean-pooling, which computes the mean response value in each grid. After pooling, the dimension of OB will be $NumObject * NumScale * 2 * 21$.

2.2 Object Bank Face (OBF)

Here we simply apply OB to face representation and name it as Object Bank Face (OBF). OBF can be seen as a high-level face representation, which is driven by models learned from the general object classes rather than the specific face classes.

Clearly, face images do not contain any general object, but face images can still form meaningful responses to OB filters. The response value reflects how likely a face patch “looks alike” the specific object in the HOG [23] feature space, as the object filters in OB do convolution on the HOG feature maps [26]. One can refer to [34] for more interesting discovery on how object detectors see this world in the HOG feature space.

3 Experimental Evaluation

In this section, we validate the OBF on face recognition task. Firstly, we introduce the evaluation dataset LFW [29] and setting of parameters. Then, we evaluate OBF on LFW and analyze the influence of different OBF setting in face recognition. Lastly, we discuss the novelty and insight of OBF.

3.1 Dataset and Setting of Parameters

In this subsection, we briefly introduce the LFW dataset, as well as the setting of parameters in our experiments.

Evaluation Dataset. LFW is a very challenging real-world face verification benchmark. The LFW dataset contains 13, 233 face images of 5, 749 persons. The face images in LFW have large variations in pose, expression, illumination, occlusion, etc. In Fig. 3, we show some example images of LFW.



Fig. 3. Example images in LFW which show the variations of face images in pose, expression, illumination and occlusion.

Face Normalization. In this work, we utilize the LFW-a images [35] provided by Lior Wolf *et al.* and crop the face image to $120 * 96$ pixels with fixed eye center locations as (29, 51) and (66, 51);

OBF Details. We use the 208 pre-trained object filters provided by Li-Jia Li *et al.* [36], which are trained on ImageNet dataset [20]. Six detection scales and 3-Level pyramid pooling are used as suggested in [1]. The dimension of the final OBF representation is $208 * 4 * 2 * 21 = 52,416$. Without additional specification, we take max-pooling as default pooling operator. In this work, whitening PCA (WPCA) is employed to reduce the dimensionality of OBF feature to 1,000, and finally cosine similarity is exploited for face matching.

3.2 Evaluation of OBF on LFW

We follow the LFW restricted protocol, which splits the LFW View 2 dataset into 10 subsets with each subset containing 300 positive pairs and 300 negative pairs. In this work, we follow the unsupervised setting. As requested by the protocol, we perform 10-fold cross-validation, among which 9 folds are used to train the WPCA model and the rest one is used for testing. The evaluation results are shown in Table. 1. Note that, we use “(sqrt)” behind a specific method to denote additional signed square root operation of features. These results are given because many previous works have shown the effectiveness of sqrt operator especially for histogram-like features.

Table 1. Comparisons between OBF and state-of-the-art methods on LFW dataset under unsupervised setting.

Method	Mean Accuracy Rate(%)
LHS [37]	73.40±0.40
LARK [38]	78.90
MRF-MLBP [39]	80.08±0.13
POEM [5]	82.71±0.59
OCLBP [40]	82.78±0.41
High-dim LBP [4]	84.08
High-dim LE [4]	84.58
I-LQP [6]	86.20±0.46
SFRD [14]	84.81
OBF	83.80±0.42
OBF (sqrt)	83.75±0.46

As can be seen from the table, the OBF surprisingly achieves 83.80% mean accuracy on LFW View 2 under unsupervised setting, which outperforms many specifically designed face representation methods such as LARK [38], POEM [5] and OCLBP [40]. The performance of OBF is even competitive to state-of-the-art methods, which seems a proof of the saying that “Everything is in the face”.

In terms of face recognition, it seems that who you are is determined by what objects are in your face.

3.3 Analysis of OBF

In this section, we analyze the role of each component in OBF. We demonstrate the sensitivity of object filters on face and the influence of different OBF setting on face recognition performance. We eventually provide a good understanding of OBF.

Sensitivity of Object Filters to Face. In order to visualize the sensitivity of object filters to face, we do statistic of the mean response maps of each object filters on face. In this experiment, we use all the 13, 233 face images from LFW dataset View 2 [29]. We randomly choose eight object filters and demonstrate the mean response maps in Fig. 4. Interesting phenomenon can be explored: big response value appears mostly around the eyes, nose and mouth. “Everything is in the face” has intuitive explanations.

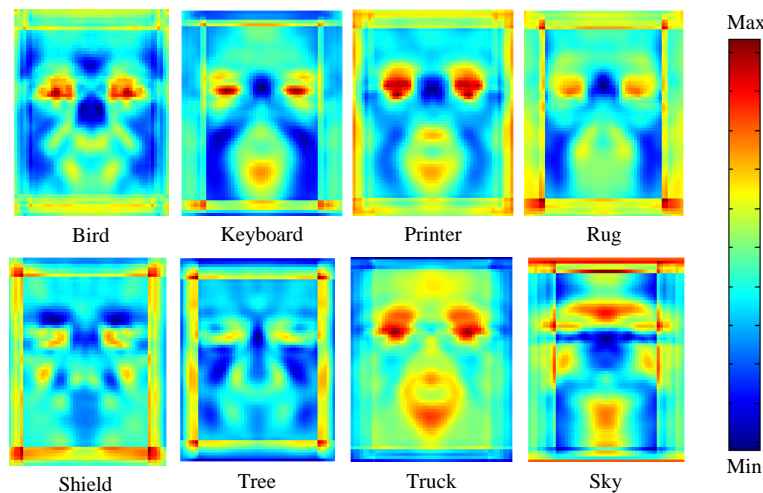


Fig. 4. Mean response maps of selected object filters on face.

Comparisons of Different Pooling Operator. As described earlier, the pooling process in OB encodes the spatial location and semantic meaning of objects in the input image. In this experiment, we analyze the performance of different pooling operator, i.e., max-pooling and mean-pooling, on face verification. We conduct the experiment on LFW View 2 under unsupervised setting. We reduce the original features to 1, 000 dimensions using WPCA and take the cosine similarity. As it is shown in Table. 2, the max-pooling outperforms mean-pooling. The same conclusion is also presented in [1], where mean-pooling outperforms max-pooling when applying OB for scene classification task.

Table 2. Comparisons of the pooling operators of OBF on LFW dataset under unsupervised setting.

Pooling Strategy	Mean Accuracy Rate(%)
OBF + Mean-Pooling	80.35±0.55
OBF + Mean-Pooling (sqrt)	80.32±0.54
OBF + Max-Pooling	83.80±0.42
OBF + Max-Pooling (sqrt)	83.75±0.46

Table 3. Comparisons of the performance of each scale of OBF on LFW dataset under unsupervised setting. The left column presents the feature map size, and the object filter kernel size is around $10 * 8$ or $8 * 10$.

Feature Map Size in Pixel	Mean Accuracy Rate(%)
S1: 61 *48	75.92±0.78
S2: 42*33	76.82±0.79
S3: 29*23	78.78±0.76
S4: 20*16	79.52±0.70
S5: 14*11	80.33±0.65
S6: 9*7	79.78±0.75

Role of Scale. In OBF, there are totally six detection scales. In this experiment, we evaluate the pooled features of each scale independently. The results are demonstrated in Table. 3. One can see that individual scale performs different with each other, and small size scales delivers better performance. Furthermore, in this experiment, we also study single object performance in different scale. The same conclusion is also hold, and the single object filter performs best on small size scale S5 or S6. This phenomenon is not hold when applying OB for scene classification task as reported in [1], where middle size scale performs better. An intuitive interpretation is that in natural image, object often occupies middle size of the input image. While for face image, it shares common visual pattern with general object in bigger face regions.

Role of Object. In this experiment, we analyze the effectiveness of different types of objects on face verification performance. It is an interesting problem that which object contributes more to the performance. The test is also taken on LFW View 2 under unsupervised setting. We reduce the dimension of original feature using WPCA and take the cosine similarity. We sort the object filters according to the 10 folder mean accuracy on LFW View 2. In Fig. 5, we report the best 20 objects and worst 20 objects together with their mean accuracies.

In this experiment, we also evaluate the performance of using different number of object filters in OBF. For computation simplicity, we conduct the experiment on LFW View 1 under unsupervised setting. The LFW View 1 has a training set of 1, 000 positive pairs and 1, 000 negative pairs, and a test set of 500 positive pairs and 500 negative pairs. We reduce the dimension of original feature using WPCA. Dimension is fixed to the minimum number of principal




Fig. 5. Performance by using OBF generated by individual object on LFW View 2 under unsupervised setting. (1) The best 20 objects; (2) The worst 20 objects.

components to preserve 80% of the training set variance. By adding one object at a time, we get the performance curve in Fig. 6. The classification accuracy increases along with the number of object filters progressively, and we believe more powerful and discriminative face representation can be constructed with the growth of number of object filters.

3.4 Discussion

Above experiments and observations suggest that general object filters can indeed achieve very promising accuracy for face recognition. This observation seems anti-intuitive, as we are not so silly to believe there are indeed trucks or printer on face. So, if we were not wrong in implementing, the effectiveness has to be hidden somewhere in the mechanism. To explain this in principal, we argue that, each object is actually composed of a large number of low-level features or mid-level attributes. In other words, an object filter models the co-occurrence of plenty of low-level filters and/or mid-level filter. Therefore, the effectiveness of OB for face recognition might essentially come from its modeling of co-occurrence of many lower-level features. We can imagine that, a truck, in a suitable scale, can be like human eye, especially in the sense of several horizontal edges.

In terms of cognitive neuroscience, as mentioned in the introduction part, it has long been controversial whether face processing is a dedicated procedure. Our findings suggest that, at least, in the low-level or even mid-level, the perception




Fig. 6. Classification performance on LFW View 1 by using OBF corresponding to accumulative objects under unsupervised setting. X axis is the number of objects. Y axis represents the average classification accuracy.

mechanism might be similar (if not the same) for face and general object processing. There is just no need to distinguish the horizontal edges on a truck or those on a face. What matters is how the low-level features are combined together in some co-occurrence mode, in the higher level of perceptual organization.

4 Combination with Low-level Features

As a high-level representation, OB is intuitively complementary for low-level feature based face representation. In the following, we firstly present the experimental setting. Then, we evaluate the proposed method on LFW.

We combine the OBF with LBP Face, SIFT Face or Gabor Face and evaluate on LFW View 2 under unsupervised setting. We reduce the dimension of original feature using WPCA and take the cosine similarity. In this experiment, we take score-level fusion with equal weight by simply summing the corresponding similarity matrices. The setting of parameters for low-level feature extraction is presented as follows.

LBP Face Details. We take the uniform LBP setting as suggested in [3], and we set the block size as $8 * 8$ and the radius as 2.

SIFT Face Details. We take the dense SIFT feature, which compute SIFT descriptor [22] on dense grid on the face image. In this experiment, we set the grid size as $4 * 4$.

Gabor Face Details. In the Gabor face, 40 Gabor wavelets with 5 scales and 8 orientations are utilized, the parameters are set as suggested in [41]. We take a $4 * 4$ mean pooling on the 40 Gabor magnitude images to reduce the feature dimensionality.

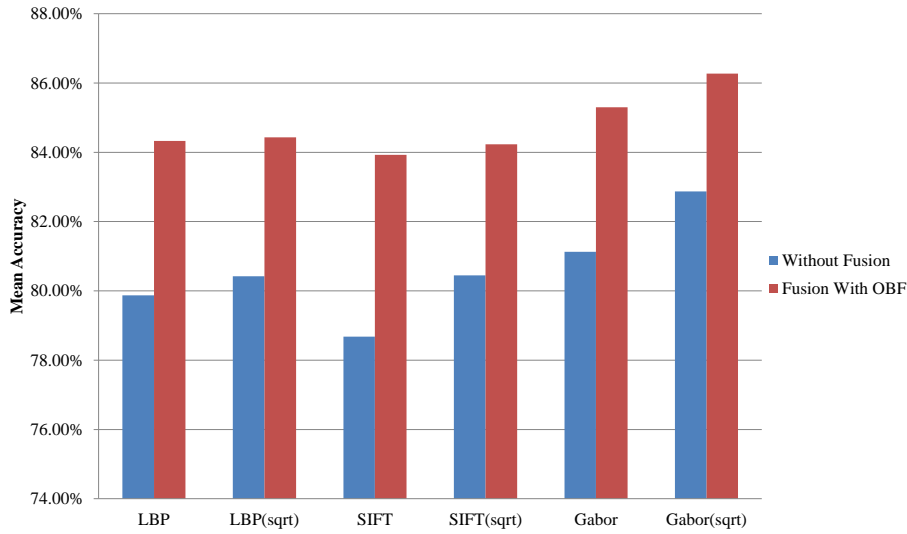


Fig. 7. Score-level fusion of OBF and low-level feature based face representation on LFW dataset under unsupervised setting.

The results are shown in Fig. 7. As it is shown, the fusing of OBF and low-level feature significantly improves the performance. Specially, we achieve 86.27% mean accuracy by fusing OBF and Gabor (sqrt), which is the state-of-the-art performance on LFW under unsupervised setting. We can conclude that OBF is an efficient complementary representation to improve face representation robustness.

5 Conclusions and Future Work

Face representation is a fundamental problem in face recognition research. In this work, we apply general object filters, i.e., OB, to face representation. The experimental evaluations on LFW demonstrate that face representation can well leverage the high-level models of general objects. Furthermore, by combining with low-level features, state-of-the-art performance can even be achieved on LFW.

The observation of this work naturally raises the conjecture whether mid-level features, e.g., attributes, for general object recognition, can better facilitate face recognition. So, in the future, we will study how general attribute can be exploited for face recognition.

Acknowledgement The work is supported by Natural Science Foundation of China under contract No. 61222211.

References

1. Li, L.J., Su, H., Lim, Y., Fei-Fei, L.: Object bank: An object-level image representation for high-level visual recognition. *International Journal of Computer Vision* (2013) 1–20
2. Liu, C., Wechsler, H.: Gabor feature based classification using the enhanced fisher linear discriminant model for face recognition. *IEEE Transactions on Image Processing* **11** (2002) 467–476
3. Ahonen, T., Hadid, A., Pietikainen, M.: Face description with local binary patterns: Application to face recognition. *IEEE Transactions on Pattern Analysis and Machine Intelligence* **28** (2006) 2037–2041
4. Chen, D., Cao, X., Wen, F., Sun, J.: Blessing of dimensionality: High-dimensional feature and its efficient compression for face verification. In: *Computer Vision and Pattern Recognition (CVPR)*, IEEE (2013) 3025–3032
5. Vu, N.S., Caplier, A.: Enhanced patterns of oriented edge magnitudes for face recognition and image matching. *IEEE Transactions on Image Processing* **21** (2012) 1352–1365
6. Hussain, S.U., Napoléon, T., Jurie, F., et al.: Face recognition using local quantized patterns. In: *British Machine Vision Conference (BMVC)*. (2012)
7. Bicego, M., Lagorio, A., Grosso, E., Tistarelli, M.: On the use of sift features for face authentication. In: *CVPRW*, IEEE (2006) 35–35
8. Albiol, A., Monzo, D., Martin, A., Sastre, J., Albiol, A.: Face recognition using hog+ebgm. *Pattern Recognition Letters* **29** (2008) 1537–1543
9. Kumar, R., Banerjee, A., Vemuri, B.C., Pfister, H.: Trainable convolution filters and their application to face recognition. *IEEE Transactions on Pattern Analysis and Machine Intelligence* **34** (2012) 1423–1436
10. Lei, Z., Yi, D., Li, S.Z.: Discriminant image filter learning for face recognition with local binary pattern like representation. In: *Computer Vision and Pattern Recognition (CVPR)*, IEEE (2012) 2512–2517
11. Cox, D., Pinto, N.: Beyond simple features: A large-scale feature search approach to unconstrained face recognition. In: *Automatic Face & Gesture Recognition and Workshops (FG)*, 2011 IEEE International Conference on, IEEE (2011) 8–15
12. Xie, S., Shan, S., Chen, X., Meng, X., Gao, W.: Learned local gabor patterns for face representation and recognition. *Signal Processing* **89** (2009) 2333–2344
13. Cao, Z., Yin, Q., Tang, X., Sun, J.: Face recognition with learning-based descriptor. In: *Computer Vision and Pattern Recognition (CVPR)*, IEEE (2010) 2707–2714
14. Cui, Z., Li, W., Xu, D., Shan, S., Chen, X.: Fusing robust face region descriptors via multiple metric learning for face recognition in the wild. In: *Computer Vision and Pattern Recognition (CVPR)*, IEEE (2013) 3554–3561
15. Kumar, N., Berg, A.C., Belhumeur, P.N., Nayar, S.K.: Attribute and simile classifiers for face verification. In: *Computer Vision and Pattern Recognition (CVPR)*, IEEE (2009) 365–372
16. Berg, T., Belhumeur, P.N.: Tom-vs-pete classifiers and identity-preserving alignment for face verification. In: *British Machine Vision Conference (BMVC)*. Volume 1. (2012) 5
17. Huang, G.B., Lee, H., Learned-Miller, E.: Learning hierarchical representations for face verification with convolutional deep belief networks. In: *Computer Vision and Pattern Recognition (CVPR)*, IEEE (2012) 2518–2525
18. Taigman, Y., Yang, M., Ranzato, M., Wolf, L.: Deepface: Closing the gap to human-level performance in face verification. In: *Computer Vision and Pattern Recognition (CVPR)*. (2014)

19. Fei-Fei, L., Fergus, R., Perona, P.: Learning generative visual models from few training examples: An incremental bayesian approach tested on 101 object categories. *Computer Vision and Image Understanding* **106** (2007) 59–70
20. Deng, J., Dong, W., Socher, R., Li, L.J., Li, K., Fei-Fei, L.: Imagenet: A large-scale hierarchical image database. In: *Computer Vision and Pattern Recognition (CVPR), IEEE* (2009) 248–255
21. Krizhevsky, A., Sutskever, I., Hinton, G.E.: Imagenet classification with deep convolutional neural networks. In: *NIPS*. Volume 1. (2012) 4
22. Lowe, D.G.: Distinctive image features from scale-invariant keypoints. *International Journal of Computer Vision* **60** (2004) 91–110
23. Dalal, N., Triggs, B.: Histograms of oriented gradients for human detection. *Computer Vision and Pattern Recognition (CVPR)* (2005)
24. Bay, H., Tuytelaars, T., Van Gool, L.: Surf: Speeded up robust features. In: *European Conference on Computer Vision (ECCV)*. Springer (2006) 404–417
25. Basura Fernando, Elisa Fromont, T.: Mining mid-level features for image classification. *International Journal of Computer Vision* **108** (2014) 186–203
26. Felzenszwalb, P.F., Girshick, R.B., McAllester, D., Ramanan, D.: Object detection with discriminatively trained part-based models. *IEEE Transactions on Pattern Analysis and Machine Intelligence* **32** (2010) 1627–1645
27. Kanwisher, N., McDermott, J., Chun, M.M.: The fusiform face area: a module in human extrastriate cortex specialized for face perception. *The Journal of Neuroscience* **17** (1997) 4302–4311
28. Grill-Spector, K., Knouf, N., Kanwisher, N.: The fusiform face area subserves face perception, not generic within-category identification. *Nature neuroscience* **7** (2004) 555–562
29. Huang, G.B., Ramesh, M., Berg, T., Learned-Miller, E.: 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)
30. Sinha, P., Balas, B., Ostrovsky, Y., Russell, R.: Face recognition by humans: Nineteen results all computer vision researchers should know about. *Proceedings of the IEEE* **94** (2006) 1948–1962
31. Grill-Spector, K., Knouf, N., Kanwisher, N.: The fusiform face area subserves face perception, not generic within-category identification. *Nature neuroscience* **7** (2004) 555–562
32. Gauthier, I., Tarr, M.J.: Becoming a greeble expert: Exploring mechanisms for face recognition. *Vision research* **37** (1997) 1673–1682
33. Gauthier, I., Anderson, A.W., Tarr, M.J., Skudlarski, P., Gore, J.C.: Levels of categorization in visual recognition studied using functional magnetic resonance imaging. *Current Biology* **7** (1997) 645–651
34. Vondrick, C., Khosla, A., Malisiewicz, T., Torralba, A.: HOGgles: Visualizing Object Detection Features. *International Conference on Computer vision (ICCV)* (2013)
35. Taigman, Y., Wolf, L., Hassner, T.: Multiple one-shots for utilizing class label information. In: *British Machine Vision Conference (BMVC)*. (2009) 1–12
36. Vision Lab, S.U.: (<http://vision.stanford.edu/projects/objectbank/>)
37. Gaurav Sharma, Sibte ul Hussain, F.J.: Local higher-order statistics (lhs) for texture categorization and facial analysis. *European Conference on Computer Vision (ECCV)* **2012** (2012)
38. Seo, H.J., Milanfar, P.: Face verification using the lark representation. *IEEE Transactions on Information Forensics and Security* **2011** (2011)

39. Arashloo, S.R., Kittler., J.: Efficient processing of mrfs for unconstrained-pose face recognition. *Biometrics: Theory, Applications and Systems* **2013** (2013)
40. Barkan, O., Weill, J., Wolf, L., Aronowitz, H.: Fast high dimensional vector multiplication face recognition. In: *International Conference on Computer vision (ICCV)*. (2013)
41. Li, Y., Shan, S., Zhang, H., Lao, S., Chen, X.: Fusing magnitude and phase features for robust face recognition. In: *Asian Conference on Computer Vision (ACCV)*. Springer (2013) 601–612