
Deep Attribute Networks

Junyoung Chung

Department of EE, KAIST
Daejeon, South Korea
jych@kaist.ac.kr

Donghoon Lee

Department of EE, KAIST
Daejeon, South Korea
iamdh@kaist.ac.kr

Youngjoo Seo

Department of EE, KAIST
Daejeon, South Korea
minerrba@kaist.ac.kr

Chang D. Yoo

Department of EE, KAIST
Daejeon, South Korea
cdyoo@ee.kaist.ac.kr

Abstract

Obtaining compact and discriminative features is one of the major challenges in many of the real-world image classification tasks such as face verification and object recognition. One possible approach is to represent input image on the basis of high-level features that carry semantic meaning which humans can understand. In this paper, a model coined deep attribute network (DAN) is proposed to address this issue. For an input image, the model outputs the attributes of the input image without performing any classification. The efficacy of the proposed model is evaluated on unconstrained face verification and real-world object recognition tasks using the LFW and the a-PASCAL datasets. We demonstrate the potential of deep learning for attribute-based classification by showing comparable results with existing state-of-the-art results. Once properly trained, the DAN is fast and does away with calculating low-level features which are maybe unreliable and computationally expensive.

1 Introduction

Many computer vision tasks such as image segmentation, object recognition, and face verification are conducted on the basis of image features providing shape, color, texture information of the image. These are generally low-level visual features that carry limited semantic meaning, and the extraction of certain low-level features such as scale-invariant feature transform (SIFT) [1], histogram of oriented gradients (HOG) [2], local binary patterns (LBP) [3] by means of signal processing is considered computationally expensive, and they are unreliable under various circumstances. To achieve high performance, various ad hoc preprocessing prior to the actual feature extraction is required. Furthermore, concatenation and combination of different low-level features enlarge the dimension of the input feature space. As an alternative, recent studies have shown that the conversion from low-level features to high-level features (henceforth referred to as attributes) that provide semantic meaning to which humans can relate has the potential to enhance the classification performance of particular computer vision tasks and can be used to describe objects unseen in the training examples [4, 5].

It was shown that humans perform almost perfectly in extremely ambiguous classification task such as unconstrained face verification [4]. Although there have been great advances recently in the understanding of the human visual cortex, humans' classification ability is hard to imitate since the human brain is still considered as a recondite area. Instead, we focus on the human descriptive ability, which is closely connected to the classification ability. For instance, one can instinctively understand that the following words, "an animal with yellow fur, a long neck, four legs, and brown

spots,” are describing a giraffe. Each word is an attribute that represents a describable visual aspect of the giraffe, and the combination of attributes can distinguish the giraffe from other animals. Based on this paradigm, a two-step framework can be considered in classification: 1) obtain an attribute descriptor (binary state or real-valued score) from an image and 2) feed it to a final classifier to perform the classification.

As an attribute-based solution, the absence or presence of various attributes in the form of a binary attribute descriptor is determined using a number of attribute classifiers. Kumar *et al.* obtained attribute score vectors from face images using attribute classifiers based on support vector machines (SVMs) on a combination of low-level features [4]. In the multi-object classification tasks, Lampert *et al.* utilized between-class attributes to overcome object classification in the case of training and test classes are disjoint [5]. Farhadi *et al.* applied both semantic and discriminative attributes to classify general objects [6].

In this paper, a supervised deep learning framework is considered to capture human-specified attributes to be used either for classification or verification. The hierarchical structure of a deep belief network (DBN) is trained to transform an input image into a form that is understandable by human without requiring the extraction of low-level features. It is known that human visual cortex exploits hierarchical feature representations of input images. For instance, the primary visual cortex V1 learns localized and oriented edge filters, and the higher-layer V2 learns linear combinations of edge filters learned in V1. Hierarchical organization in visual cortex enables the human brain to understand the scenes. In general, DBNs are pre-trained in an unsupervised manner, and the activations of the topmost layer are abstract and compact binary descriptor which is frequently used as input of classifiers such as SVMs. DBNs have shown promising performance in many applications such as phone recognition [7], object classification [8], and face verification [9]. However, the binary coded descriptors provided by DBNs are not understandable by humans. If one could train the DBN in a supervised manner to output attribute descriptor which can be interpreted by humans, it would make for an interesting study. We are motivated by the ability of the DBN that transforms the input images into abstract representations, and this nature fits well with our framework: modeling the human visual cortex to obtain the human description (attributes). The advantage of the proposed DBN-based model is that it is fast once trained and does not require the extraction of low-level features from unknown images.

The rest of this paper is organized as follows. We first present the background of the DBN in Section 2. The details of the proposed model are described in Section 3. Experiments on face and object datasets are reported in Section 4, followed by a conclusion in Section 5.

2 Deep belief networks

The DBN is composed of many hidden layers, the activations of the previous layer are the inputs of the subsequent layer, and each layer tends to learn hierarchies of feature representation of the input data. The hierarchical structure of the DBN encourages the model to gain representational power that can represent highly non-linear and highly varying functions [10]. However, it was known difficult to train directed belief networks with deep structure due to the well-known phenomenon of explaining away before Hinton *et al.* proposed an efficient algorithm applying restricted Boltzmann machines (RBMs) [11] to train a DBN in a greedy layer-wise manner [12]. After greedy layer-wise training, the resulting model has bipartite connections at the top two layers that form an RBM, and the remaining layers are directly connected [13]. The following sections will briefly review the background information of the DBN and its building block, the RBM, before introducing our model.

2.1 Restricted Boltzmann machines and Gaussian units

The RBM is an undirected bipartite graph with a visible layer \mathbf{v} , which represents the input data, is connected to a hidden layer \mathbf{h} , which captures the underlying structures of the \mathbf{v} using symmetric weighted connections. The joint distribution of the RBM can be formulated as follows:

$$P(\mathbf{v}, \mathbf{h}; \theta) = \frac{1}{Z(\theta)} \exp(-E(\mathbf{v}, \mathbf{h}; \theta)), \quad (1)$$

where $\theta = \{W, b, c\}$ is the model parameter set of the RBM and $Z(\theta)$ is the partition function. $E(\cdot)$ denotes the energy function and is defined as follows:

$$E(\mathbf{v}, \mathbf{h}; \theta) = - \sum_{i=1}^V \sum_{j=1}^H v_i W_{ij} h_j - \sum_{i=1}^V c_i v_i - \sum_{j=1}^H b_j h_j, \quad (2)$$

where W_{ij} is the symmetric connection between a visible unit v_i and a hidden unit h_j , while c_i and b_j are biases of each unit respectively. V and H are the number of units in the visible and the hidden layers.

In general, RBMs use Bernoulli units for both of the visible and the hidden layers. When the input data are not binary-valued, RBMs reach limitation with modeling real-valued data. Instead of using general RBMs, we should consider using modified RBMs that replace the Bernoulli units in the visible layer by linear units with Gaussian noise. Hence, the modified RBMs are called Gaussian-Bernoulli RBMs [14] and the energy function is defined as follows:

$$E(\mathbf{v}, \mathbf{h}; \theta) = - \sum_{i=1}^V \frac{(v_i - c_i)^2}{2\sigma_i^2} - \sum_{i=1}^V \sum_{j=1}^H \frac{v_i}{\sigma_i} W_{ij} h_j - \sum_{j=1}^H b_j h_j, \quad (3)$$

where σ_i denotes the variance of the i^{th} visible unit. Since there are no intra-layer connections, units in the hidden layer are conditionally independent given the visible layer, and vice-versa. This property makes the inference easy. The conditional distributions of Gaussian-Bernoulli RBMs can be computed efficiently using the following equations:

$$P(h_j = 1 | \mathbf{v}; \theta) = f\left(\frac{1}{\sigma_i} \sum_{i=1}^V v_i W_{ij} + b_j\right),$$

$$P(v_i | \mathbf{h}; \theta) = \mathcal{N}(v_i | \sigma_i \sum_{j=1}^H h_j W_{ji} + c_i, \sigma_i^2), \quad (4)$$

where $f(\cdot)$ denotes the sigmoid function and $\mathcal{N}(\cdot | \mu, \sigma^2)$ is a Gaussian distribution. RBMs are trained with maximum-likelihood learning (or minimize negative log-likelihood); however, computing the exact gradient of the log-likelihood is intractable. Therefore, we use alternate method such as contrastive divergence (CD), to approximate the gradient of the log-likelihood.

2.2 Supervised learning of RBMs

Typically, the activations of the topmost layer of DBNs are fed into classifiers such as SVMs or softmax regression model to perform classifications. We have mainly described RBMs as commonly-used building blocks of DBNs; however, recent studies address the possibility of RBMs as classifiers themselves, not as feature extractors for other algorithms. Larochelle and Bengio [15] showed impressive results by training an RBM jointly with its input data and labels. In [15], softmax regression model is combined with the RBM, modifying the energy function and the inference rule. In our approach, the DBN is used to predict the attribute scores of the input data. In order to meet our goal, we changed the RBM to gain the discriminative ability to output a vector prediction. The canonical way to make a DBN discriminative is fine-tuning which slightly changes the parameters to gain better generalization. From our perspective, the supervised learning of the RBM fits better than fine-tuning and this will be discussed in the following sections.

3 Algorithms

Obtaining attribute scores from an input image is equivalent to modeling the procedure of human brains, understanding semantic meaning from the visual information processed by the visual cortex. Existing attribute-based approaches used complex combinations of low-level features and fed them into SVMs to obtain attribute scores from the input data [4, 5, 6]. Our model is based on the DBN, which is an appropriate model to describe the human visual cortex and automatically learns feature representations of the input images. This property is advantageous when adopting our model to many other applications in computer vision. Although the hyper parameters of the model might change, it still learns the underlying structures of the input images automatically.

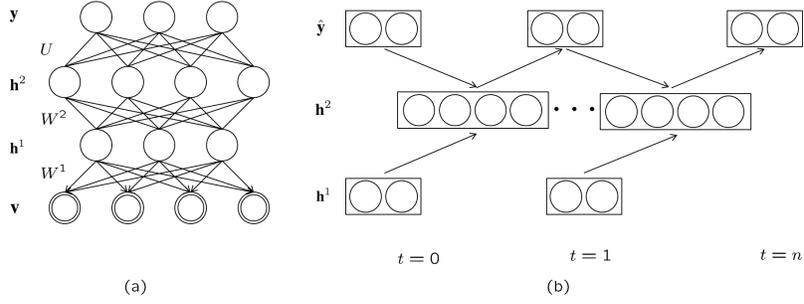


Figure 1: (a) shows the graphical model of the DAN composed with multiple layers including the attribute layer y , (b) presents our inference procedure, running mean-field approximation for n iterations, keeping h^1 fixed.

3.1 Modeling the human visual cortex

According to the neuroscience, the neurons in primary visual cortex V1, function as edge detectors, and the neurons in subsequent layer V2 are likely to learn contour and corner detectors which are combination of the edge detectors. These neurons can represent an input image into a sparse representation which is known to be advantageous in some reasons. Features are likely to be more separable in high-dimensional spaces, sparse constraint let the training procedure become efficient and the classification become robust to noise [16].

3.1.1 Introducing sparsity to RBMs

In [17], it was shown that introducing sparse regularization term to the objective function encourages the model to learn sparse representations of the input images. Sparse variant of the DBN has successfully learned sparse representations of the input images. Given a training set $\mathcal{X} = \{x^{(1)}, \dots, x^{(m)}\}$, the objective function of sparse RBMs (SRBMs) can be written as follows:

$$\text{minimize}_{\theta} - \sum_{k=1}^m \log \sum_{\mathbf{h}} P(\mathbf{v}^{(k)}, \mathbf{h}^{(k)}) + \frac{\lambda}{2} \sum_{i=1}^V \sum_{j=1}^H \|W_{ij}\|^2 + \beta \sum_{j=1}^H \left(\rho - \frac{1}{m} \sum_{k=1}^m \mathbb{E}[h_j^{(k)} | \mathbf{v}^{(k)}] \right)^2, \quad (5)$$

where the second term is regularization of weights, ρ is the target sparsity, and β is a constant that controls the weight of the sparse regularization term ($\rho^1 = 0.1$, $\rho^2 = 0.2$, and $\beta = 2$ sufficed in our experiments).

SRBM is an appropriate method to learn edge filters that can be considered as “biologically-inspired” representations. Stacking SRBMs forms a DBN, and contour and corner filters are learned as the second layer bases. Furthermore, we propose deep attribute networks (DANs) to predict the score vector of k -binary units that correspond to k attributes. Figure 1 depicts the schematic view of our model. Given a training set $\mathcal{X} = \{x^{(1)}, \dots, x^{(m)}\}$ and labels $\mathcal{Y} \in \{0, 1\}^m$, the joint probability of the DAN is defined as follows:

$$P(\mathbf{v}, \mathbf{h}^1, \mathbf{h}^2, \mathbf{y}; \theta) = P(\mathbf{v} | \mathbf{h}^1; \theta) P(\mathbf{h}^1, \mathbf{h}^2, \mathbf{y}; \theta), \quad (6)$$

where $\theta = \{W, b^1, b^2, U, d\}$ is the model parameter set of the SRBM. Since the connections of the bottom layers are directed, the first block (involving \mathbf{v} and \mathbf{h}^1) is trained by the SRBM. The second block (involving \mathbf{h}^1 , \mathbf{h}^2 , and \mathbf{y}) is trained by discriminative SRBM (DSRBM) in a supervised manner. The energy function of the DSRBM is given as follow:

$$E(\mathbf{h}^1, \mathbf{h}^2, \mathbf{y}; \theta) = - \sum_{i=1}^{H_1} \sum_{j=1}^{H_2} h_i^1 W_{ij}^2 h_j^2 - \sum_{i=1}^{H_1} b_i^1 h_i^1 - \sum_{j=1}^{H_2} b_j^2 h_j^2 - \sum_{j=1}^{H_2} \sum_{k=1}^Y h_j^2 U_{jk} y_k - \sum_{k=1}^Y d_k y_k, \quad (7)$$

where H_1 and H_2 are the number of units in the first hidden and the second hidden layers. U_{jk} is the connection between a hidden unit h_j^2 and an attribute unit y_k , and d_k is the bias term of each

attribute unit, and the rest of the variables are same as in Equation 2. Inference is straightforward: since there are no intra-layer connections between attribute units, conditional distributions become independent given \mathbf{h}^2 . The conditional distributions are formulated as follows:

$$P(h_i^1 = 1 | \mathbf{h}^2; \theta) = f\left(\sum_{j=1}^{H_2} h_j^2 W_{ij}^2 + b_i^1\right), \quad (8)$$

$$P(h_j^2 = 1 | \mathbf{h}^1, \mathbf{y}; \theta) = f\left(\sum_{i=1}^{H_1} h_i^1 W_{ij}^2 + \sum_{k=1}^{\mathcal{Y}} y_k U_{jk} + b_j^2\right), \quad (9)$$

$$P(y_k | \mathbf{h}^2; \theta) = f\left(\sum_{j=1}^{H_2} h_j^2 U_{jk} + d_k\right). \quad (10)$$

In contrast to [15], we do not combine softmax regression model with the SRBM. Instead, we put an additional attribute layer with k -binary units to output a vector prediction of the attribute scores. A stochastic gradient descent, such as CD, is used to update the parameters as follows:

$$\Delta\theta = \epsilon \left(- \left\langle \frac{\partial E(\mathbf{h}^1, \mathbf{h}^2, \mathbf{y}; \theta)}{\partial \theta} \right\rangle_{\text{data}} + \left\langle \frac{\partial E(\mathbf{h}^1, \mathbf{h}^2, \mathbf{y}; \theta)}{\partial \theta} \right\rangle_{\text{reconstruction}} \right), \quad (11)$$

where $\theta = \{W^2, b^1, b^2, U, d\}$ is the model parameter set and ϵ is a learning rate ($\epsilon^1 = 0.001$ and $\epsilon^2 = 0.01$ sufficed in our experiments). The minus sign and the plus sign are reversed, because the optimization problem is minimizing the negative log-likelihood.

3.2 Prediction

Given a test image, borrowing ideas from the training procedure in deep Boltzmann machines (DBMs), we are able to predict the attribute score vector $\hat{\mathbf{y}}$ by mean-field approximation [13]. Our procedure is depicted in Figure 1. We fixed \mathbf{h}^1 and ran 10 iterations to obtain $P(\hat{\mathbf{y}} | \mathbf{h}^2)$. In our approach, we used the energy of each unit in $\hat{\mathbf{y}}$ as a score value for each attribute.

4 Experiments

Two experiments on real world images are considered to evaluate the performance of the DAN. Firstly, we conducted experiments for unconstrained face verification task on the labeled faces in the wild (LFW) dataset [18]. We compared the verification performance between the descriptor-based model using low-level features and attribute-based model which is the proposed model of this paper. Secondly, we performed experiments on real-world object recognition on the a-PASCAL dataset [6]. In object recognition, we compared the classification performance with existing attribute-based models using SVM attribute classifiers and proposed model using DANs for attribute classification. Many of the hyper parameter values were determined by performing tests on the validation set. The details are described in the following subsections.

4.1 Unconstrained face verification

The LFW dataset is organized into two views: a development set consists of 2,200 training pairs and 1,000 validation pairs; randomly generated 10-fold set of 6,000 pairs to evaluate final verification performance. We used View 1 to determine the parameters of the proposed model, and report the mean accuracy of 10-fold cross validation results using the image restricted configuration (extra pairs are not generated).

a-LFW: We have developed new annotations on the LFW, for exploring the attribute-based face verification. The number of collected annotations is 36 per face image, which can be judged objectively and inferred from the appearance of the faces. The labeling task becomes ambiguous for certain images, so that majority voting was used on 9 labels provided by 9 individuals participating in the experiment.

Table 1: Attribute classification accuracies of the DAN, the HOG and the LBP. AUCs are measured for all attributes and averaged.

Method	DAN	HOG	LBP
avg-AUC	0.86	0.82	0.84

Face representation: We detected 9 facial landmarks following the method described in [19]. Each face image was aligned using 2-D affine transform based on the detected landmarks, and 10 patches were generated: 9 patches were cropped by counting each facial landmark as a center point, and 1 patch was resized square image of the global face (each patch is 20×20). The generated patches were normalized into zero-mean and unit variance, and ZCA whitening was performed to eliminate the second-order correlations between adjacent pixels.

We also prepared two low-level descriptors to compare the face verification performance with the high-level descriptor obtained by the DAN: LBP and HOG. For the LBP, we used 8 uniform-spaced circular neighbor sets with radius 3 and 59-code encoding, and for the HOG, 32 bins were used for quantization. Both features were extracted from the training set (generated patches), which were used to train the DAN (without performing normalization and ZCA whitening). Extracted low-level features were used as input features of SVM classifier with an RBF kernel to build attribute classifiers. The attribute score vector extracted by SVM based attribute classifiers was input features of the final face verification classifier which is also an SVM with the RBF kernel.

Attribute classification & face verification: Generated patches were stacked into a batch file and were used as the input data of the first block of the DAN which is a Gaussian SRBM (GSRBM) in order to learn edge filters (First layer bases). The number of the units in the visible layer and in the first hidden layer was set to 400. After the first block was trained, we froze the parameters $\theta^1 = \{W^1, b^1, c\}$, and computed the expectations of the first hidden activations $\mathbb{E}[\mathbf{h}^1 | \mathbf{v}]$, given 10 patches cropped from a same face image. We concatenated the computed activations into an array, and repeated this work for every training sample. The second block, which is a DSRBM, used the concatenated arrays (400×10 , [1, 4000]) and the ground truth labels of the attributes ([1,36]) as its input. The number of the units of the first hidden layer and the second hidden layer were set to 4000 and 3200, respectively. Test images were processed in the same method described above, and the attribute scores were obtained by using the method described in Section 3.2. Our final classifier was the SVM with radial basis function (RBF) kernel trained using LIBSVM [20]. The learned first layer bases are depicted in Figure 2.

Experimental results on attribute classification are shown in Table 1. We applied area under ROC curve (AUC) to measure the classification accuracy of each attribute. The results show that our method outperforms the LBP and the HOG. AUCs of attributes classified by the DAN are shown in Table 2, and the face verification results are shown in Table 3 and the ROC curves of each method are drawn in Figure 3. The experimental results demonstrate that the attribute-based face verification consistently outperforms the low-level feature based face verification. Since we used less attributes than [4], our performance is slightly lower than their work. We pose that the authors of [4] used 65 attributes while we used only 36 attributes. The reason why we did not use same attributes like as [4] is due to the copyright problem that we couldn't get access to the training data which were used in the previous work. The verification performance of the ground truth labels of the attributes was $93.4 \pm 0.97\%$, and the verification performance of humans was reported in [4] which is 99.2% .

4.2 Object recognition

In real-world object recognition task, we evaluate our proposed model on the a-PASCAL dataset, which is developed for attribute-based classification.

a-PASCAL: Farhadi *et al.* collected 6340 training images, and 6355 test images containing 20 object classes from the PASCAL VOC 2008 challenge dataset [6]. Each image is annotated with 64 attributes. The examples of the annotated attributes are 2-D boxy, tail, wheel, side mirror, window, shiny, etc.. We exploited 64 attributes as a descriptor to classify the objects.

Table 2: Attribute classification accuracies of the DAN: AUCs of 36 attributes are represented.

Attribute	AUC	Attribute	AUC	Attribute	AUC
Male	0.96	Gray hair	0.83	Short hair	0.76
Female	0.97	Thick eyebrows	0.81	Curly hair	0.74
Mustache	0.93	No eye-wear	0.94	Eyes slightly opened	0.74
Beard	0.94	Wearing glasses	0.96	Eyes opened	0.82
Heavy makeup	0.95	Wearing sunglasses	0.98	Closed eyes	0.81
Wearing lipsticks	0.96	Wearing a hat	0.88	Side lip wrinkles	0.87
Smiling	0.95	Small eyes	0.80	Mouth closed	0.89
Bald	0.82	Visible forehead	0.76	Athlete	0.90
Receding hairline	0.72	Old	0.88	Caucasian	0.85
All back hair	0.68	Mouth slightly opened	0.77	Asian	0.94
Blond	0.84	Mouth widely opened	0.87	African-American	0.96
Black hair	0.81	Teeth invisible	0.92	Bangs	0.82

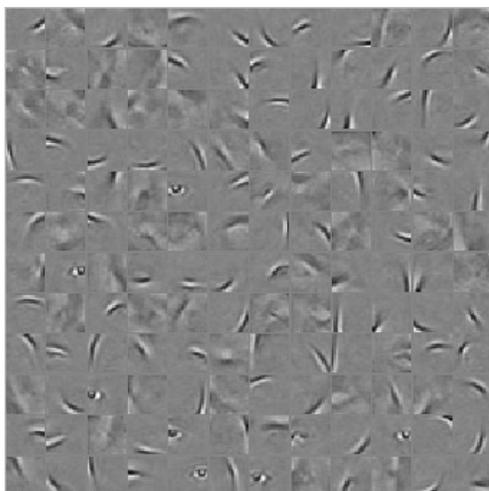


Figure 2: A subset of the 1st layer bases learned by GSRBM on 20×20 patches generated from the LFW. The GSRBM has 400 units for both of the visible layer and the hidden layer.

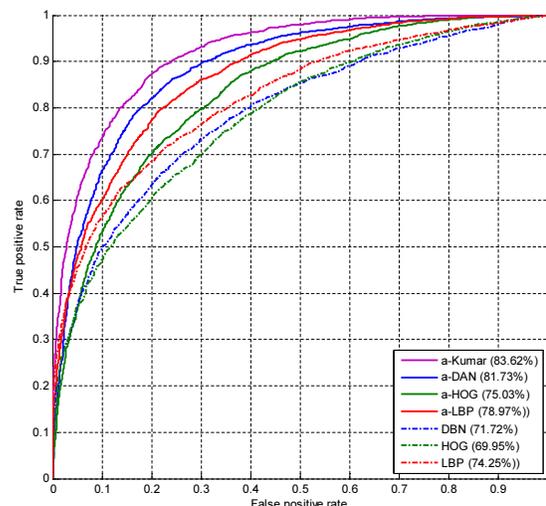


Figure 3: Face verification results on the LFW dataset.

For the fair comparison between algorithms, the developers of the dataset [6] provide bag-of word style features which are commonly used in the object recognition task ¹ and we also used these features instead of using raw pixel values of the object images.

Attribute classification & Object recognition: The base features that we used are based on histogram, so that the data have sparseness; therefore, we used the SRBM to train the first block. We simply made an array for each sample by concatenating the features (The dimension of each array was [1, 1392]). We linearly interpolated each element into a range [0, 5], and labeled 1 if the value was greater than 1. The dimension of the visible layer and the first hidden layer of the SRBM were 1392 and 800, respectively. The second block was trained by DSRBM taking the computed activations of the first hidden layer (800×10 , [1, 8000]) and annotated attributes ([1, 64]) as its input and the number of the second hidden layer was set to 8000 .

Table 4 shows the object recognition results. The a-PASCAL dataset is heavily biased to the “person category” (5071 of 12695), which implies that the overall accuracy on this task is dominated by the “person category”. Therefore, “mean per class” accuracy which represents the average classifica-

¹<http://vision.cs.uiuc.edu/attributes/>

Table 3: Face verification results of attribute-based verification and direct verification.

Descriptor	Att-based verification	Direct verification
DAN	$81.73 \pm 0.97\%$	$71.72 \pm 2.68 \%$
LBP	$75.03 \pm 2.88\%$	$74.25 \pm 1.3 \%$
HOG	$78.97 \pm 1.92\%$	$69.95 \pm 1.79 \%$
Kumar <i>et al.</i> [4]	$83.62 \pm 1.58\%$	-

Table 4: Classification accuracies on the a-PASCAL dataset.

Method	Overall	Mean per class
DAN	53.7%	56.1%
Latent SVM + Δ_{new} [21]	59.1%	50.8%
Latent SVM + $\Delta_{0/1}$ [21]	62.1%	46.2%
Semantic Att + SVM [6]	54.6%	28.4%
Base features +SVM [6]	58.5%	35.5%

tion accuracy of the entire classes (mean of classification accuracies of 20 objects) is more reliable measure than overall accuracy. Since our model is not biased to the “person category”, we scored 56.1% on the class mean accuracy and 53.7% on the overall accuracy.

5 Conclusion

This paper is motivated by recent studies using visual attributes to solve challenging classification tasks in computer vision. In particular, it may take a lot of efforts to label the attributes of the training data. However, when the variety of the object classes is very large, the idea of using labeled attributes in supervised learning seems to be much more reasonable and even if an unseen object was never observed in the training data, we still have a chance to predict attribute information of the object. Attribute-based method is closely related with image understanding and semantic image retrieval system. It is common to train an attribute classifier with hand-crafted low-level features, which are computationally expensive. We tackled this problem by introducing deep learning based attribute classification. The proposed model DAN can automatically learn the hierarchical feature representations from the input data, reflecting the attribute information of the training set. We demonstrate that the DAN can effectively learn attribute scores from target objects, including human faces. The contribution of this paper is to demonstrate the potential of deep learning in attribute-based classification on unconstrained face verification and real-world object recognition tasks. To the best of our knowledge, our work is the first attempt to apply deep learning in attribute-based classification based on a number of human-specified attributes. Since the deep learning based approach does not require the extraction of hand-crafted low-level features, our study holds promise for applying deep learning to attribute-based classification. In our future work, we consider the use of convolutional DBN which is known to be translation invariant and scalable to large images.

References

- [1] Lowe, D. (2004) Distinctive image features from scale-invariant keypoints. *IJCV*, **60**(2), 91–110.
- [2] Dalal, N. and Triggs, B. (2005) Histograms of oriented gradients for human detection. In *CVPR* Vol. 1, pp. 886–893.
- [3] Ahonen, T., Hadid, A., and Pietikainen, M. (2006) Face description with local binary patterns: Application to face recognition. *PAMI*, **28**(12), 2037–2041.
- [4] Kumar, N., Berg, A., Belhumeur, P., and Nayar, S. (2009) Attribute and simile classifiers for face verification. In *CVPR* pp. 365–372.
- [5] Lampert, C., Nickisch, H., and Harmeling, S. (2009) Learning to detect unseen object classes by between-class attribute transfer. In *CVPR* pp. 951–958.
- [6] Farhadi, A., Endres, I., Hoiem, D., and Forsyth, D. (2009) Describing objects by their attributes. In *CVPR* pp. 1778–1785.
- [7] Mohamed, A., Dahl, G., and Hinton, G. (2009) Deep belief networks for phone recognition. In *NIPS Workshop on Deep Learning for Speech Recognition and Related Applications*.
- [8] Lee, H., Grosse, R., Ranganath, R., and Ng, A. (2009) Convolutional deep belief networks for scalable unsupervised learning of hierarchical representations. In *Proceedings of the 26th Annual International Conference on Machine Learning ACM* pp. 609–616.
- [9] Nair, V. and Hinton, G. (2010) Rectified linear units improve restricted boltzmann machines. In *Proc. 27th International Conference on Machine Learning*.
- [10] Le Roux, N. and Bengio, Y. (2008) Representational power of restricted boltzmann machines and deep belief networks. *Neural Computation*, **20**(6), 1631–1649.
- [11] Smolensky, P. (1987) Information Processing in Dynamical Systems: Foundations of Harmony Theory. In Rumelhart, D. E., McClelland, J. L., et al., (eds.), *Parallel Distributed Processing: Volume 1: Foundations*, pp. 194–281 MIT Press Cambridge.
- [12] Hinton, G., Osindero, S., and Teh, Y. (2006) A fast learning algorithm for deep belief nets. *Neural computation*, **18**(7), 1527–1554.
- [13] Salakhutdinov, R. and Hinton, G. (2009) Deep boltzmann machines. In *Proceedings of the international conference on artificial intelligence and statistics*.
- [14] Welling, M., Rosen-Zvi, M., and Hinton, G. (2005) Exponential family harmoniums with an application to information retrieval. *NIPS*, **17**, 1481–1488.
- [15] Larochelle, H. and Bengio, Y. (2008) Classification using discriminative restricted boltzmann machines. In *ICML* pp. 536–543.
- [16] MarcAurelio Ranzato, Y., Boureau, L., and LeCun, Y. (2007) Sparse feature learning for deep belief networks. *NIPS*, **20**, 1185–1192.
- [17] Lee, H., Ekanadham, C., and Ng, A. (2008) Sparse deep belief net model for visual area V2. *NIPS*, **20**, 873–880.
- [18] Huang, G. B., Ramesh, M., Berg, T., and Learned-Miller, E., Labeled Faces in the Wild: A Database for Studying Face Recognition in Unconstrained Environments. Technical Report 07-49, University of Massachusetts, Amherst (October, 2007).
- [19] Guillaumin, M., Verbeek, J., and Schmid, C. (2009) Is that you? Metric learning approaches for face identification. In *ICCV Ieee* pp. 498–505.
- [20] Chang, C. and Lin, C. (2011) LIBSVM: a library for support vector machines. *ACM Transactions on Intelligent Systems and Technology (TIST)*, **2**(3), 27.
- [21] Wang, Y. and Mori, G. (2010) A discriminative latent model of object classes and attributes. *Computer Vision–ECCV 2010*, pp. 155–168.