# Multi-label Object Attribute Classification using a Convolutional Neural Network

Soubarna Banik, Mikko Lauri, Simone Frintrop {4banik, lauri, frintrop}@informatik.uni-hamburg.de

Department of Informatics, Universität Hamburg

**Abstract.** Objects of different classes can be described using a limited number of attributes such as color, shape, pattern, and texture. Learning to detect object attributes instead of only detecting objects can be helpful in dealing with a priori unknown objects. With this inspiration, a deep convolutional neural network for low-level object attribute classification, called the Deep Attribute Network (DAN), is proposed. Since object features are implicitly learned by object recognition networks, one such existing network is modified and fine-tuned for developing DAN. The performance of DAN is evaluated on the ImageNet Attribute and a-Pascal datasets. Experiments show that in comparison with state-of-the-art methods, the proposed model achieves better results.

## 1 Introduction

In recent years, the problem of attribute classification has attracted substantial interest within the computer vision community. Attributes can be viewed as descriptive properties of objects [5]. Different types of object attributes have been explored in research: low-level visual adjectives (such as *color*, *shape*) [16], inherent object characteristic (such as *material*) [5] or high-level object components (such as has-tail, wears-sunglasses) [2,5]. Recognizing object attributes has useful applications. Objects, though belonging to different classes, share low-level attributes such as color, shape, material. Attributes help to compare objects and categorize them. Unknown objects can be described by means of a few reference attributes. Their similarity/dissimilarity with known objects may help to infer some of their other characteristics. For example, a robot which is familiar with an orange and knows how to grasp it, can deduce how to do the same for an unknown but similarly shaped object, if it learns the shape attribute round. Another important application of attribute learning is attribute-level object localization. When we do not have exact information about an object's class but only have partial knowledge about it, such as its shape or color, an attribute-level search can help in localizing the desired target with the help of the descriptive attributes. For known objects, the additional attribute information can help in further reducing the search space. Only attributes which can be described using words, i.e., semantic attributes [16], are considered in this work. Traditional approaches [3,8,16] use handcrafted features for low-level attribute classification. In a few recent approaches, learned features are used instead of fixed handcrafted



Fig. 1: Overview of the proposed multi-label attribute classification system.

features [2,20]. However, the types of attributes learned are limited to complex high-level attributes like human attributes (gender, dress type, facial expression, wears-hat) or object-part attributes (wheel, side-mirror, window). These attributes are very specific to certain objects and are not universally applicable. Learning generic low-level attributes like color, shape, pattern, and texture can be beneficial as they can describe a wide range of objects. For this paper, we consider learning low-level object attributes.

Object attributes are often correlated. For example, the attribute *mustache* usually co-exists with the attribute *male*. This correlation is treated in different ways in research. In [16] [23], the correlation between attributes is ignored and the attributes are learned independently of each other. On the other hand, in [20] [9], attributes are jointly learned by following a multi-task learning method and by explicitly modeling the correlation between the attributes. Prediction of one attribute will increase the probability of the correlated attributes in this case. However, low-level attributes like colors and shapes are independent of each other. Colors such as *green* and *brown* may co-exist heavily in vegetation related images, but this correlation is not universal and may not exist in other types of images. Embedding this correlation into the learning process will make the model more prone to overfitting the training data.

In this paper we propose a model, named the *Deep Attribute Network* (DAN) that learns low-level object attributes belonging to the attribute groups *color*, *shape*, *pattern* and *texture*. For this purpose, an existing object classification convolutional network (ConvNet) is adapted. ConvNet features have been used for attribute classification in the past [3,14]. However, the features are used without any customization with respect to the target task. These features, extracted from the higher layers of the ConvNet are more specific to the objects, rather than to the object attributes. In our approach, the layers of the base object network are finetuned for the task of object attribute classification. Object attribute classifications are incorporated into the base network. We learn the attributes jointly so that the model can share low-level features like edges and corners across classes. The class label correlations are not enforced explicitly during the learning process. Fig. 1 shows an overview of the proposed attribute classification system. To the

best of our knowledge, ours is the first deep learning model that learns low-level object attributes and uses a multi-label classification approach.

Our network is trained on the ImageNet Attribute dataset [16]. For evaluation we use the ImageNet Attribute dataset [16] as well as the a-Pascal dataset [5]. We compare the performance of our model with that of state-of-the-art methods and show that the proposed model outperforms the existing approaches. The rest of the paper is organized as follows: section 2 discusses the state-of-the-art in attribute classification. The architecture of the proposed model is described in section 3. Section 4 discusses the performance of our model in comparison with state-of-the-art methods. We conclude and present the future work in section 5.

# 2 State of the art

Traditional attribute learning methods [5,6,12,16] treat the multi-label object attribute classification problem as a set of n independent binary classification problems, where n is the number of target attribute classes. These approaches use handcrafted feature descriptors, such as color and texture descriptors [16]. In [3,13,14], generic feature descriptors are generated from object classification ConvNets, and then the same binary classification approach is followed.

In a few object recognition tasks [5,12], attributes are used as an intermediate representation. Farhadi et al. [5] shift the goal of object recognition from naming the objects to describing the objects, which helps in case of unseen object classes or classes with very few examples. Lampert et al. [12] take a similar approach and learn the relationships between image and attributes, and between attributes and object classes. This allows the transfer of information from training classes to test classes through the intermediate attribute representation. With proper choice of attributes, their method can detect new object categories without any re-training step.

Most of the early studies on attribute classification do not consider the correlations among attributes and train each attribute classifier independently [5,16]. Recent methods [9,20] preserve the contextual correlation between attributes by using mechanisms such as contextual Restricted Boltzmann Machines (RBM) or special graph based formulations. In other related multi-label classification problems such as image captioning and multi-class multi-label object classification, this correlation is learned using a unified ConvNet-Recurrent Neural Network based framework [19].

In addition to this, there exist a few approaches, where a joint learning technique is followed with a focus on de-correlating the attributes [1,10]. These methods encourage sharing more information among attributes belonging to the same group and less information for attributes from different groups, thereby promoting certain kinds of correlation over others. Jayaraman et al. [10] implement this by weighting the feature channels differently for different attribute groups, although in a non-deep learning set up. Abdulnabi et al. [1] propose n ConvNets (for n target attributes), that share a latent fully-connected layer before the output layer. They employ regularization techniques during the training process for

#### 4 Soubarna Banik, Mikko Lauri, Simone Frintrop

intra-group feature sharing and inter-group feature competition. The provision of having an individual ConvNet for each attribute requires a large number of parameters. Also, Abdulnabi et al. [1] follow a late fusion approach (just before the output layer), which generates redundant information in the early layers of the individual networks.

Our goal is to formulate the multi-label classification problem for low-level object attributes using a single ConvNet, where the attributes are learned jointly. The nodes in the last fully-connected layer of a ConvNet share the hidden layers of the network. This enables them to capture the relationship between the classes implicitly, and hence no additional mechanism is incorporated to enforce the class correlations.

# 3 Architecture

ConvNets trained on images learn simple features like Gabor filters or color blobs in the first few layers, irrespective of the training objectives [21]. These features are useful for learning our target attribute classes - *color*, *shape*, *pattern* and *texture*. Features computed in the last layers of ConvNets are more specific to the task and the dataset at hand [21]. Object classes often contain instances of varied colors (e.g. white tigers, yellow tigers), patterns (e.g. dog with/without spots) or other variable features. This encourages the object recognition network to ignore these attributes in the later layers in order to generalize across the object class instances. Based on these observations, we finetune the last few layers of an object classification ConvNet (VGG-16 [17]) in order to adapt to the attribute classification task.



Fig. 2: Architecture of the Deep Attribute Network (DAN). The network is adapted from the VGG-16 model [17], by replacing the last two fully-connected layers with the new layers fcA and fcB. DAN is finetuned till layer  $conv5_1$ .

The proposed Deep Attribute Network (DAN) for low-level attribute classification is shown in Fig. 2. The network is trained using transfer learning starting from the VGG-16 network, pre-trained on the ImageNet dataset [15]. All convolutional layers and the first fully-connected layer (fc6) of VGG-16 are retained in DAN. Two new fully-connected layers fcA with 1024 units and fcB with 25 units are added at the end. The 25 units in the fcB layer correspond to the 25 target attribute classes in the ImageNet Attribute dataset [16] belonging to the attribute groups *color*, *shape*, *pattern* and *texture*. All convolutional layers and fully-connected layers except fcB use ReLU activation.

Each image in the ImageNet dataset [15] represents a single object class. Hence, VGG-16 uses softmax as the final layer to represent a categorical distribution over the mutually exclusive object classes. Object attributes, however, are not mutually exclusive and multiple attributes can co-occur in an object. For example, the Swedish flag is *blue*, *yellow*, and *rectangular* at the same time. Therefore, the softmax layer of VGG-16 is replaced by a sigmoid layer in DAN. The last fully-connected layer *fcB* and the sigmoid layer are represented jointly in Fig. 2. The sigmoid layer computes the sigmoid score  $\sigma(o_j) = \frac{1}{1+e^{-o_j}}$  for each element  $o_j$  in layer *fcB*, thereby converting the classifier score for each attribute class to a probability score in the range [0, 1].

**Training:** DAN is trained by minimizing a weighted cross-entropy loss using the stochastic gradient descent (SGD) algorithm. Cross-entropy is one of the popular choices as a cost function as it is non-negative, and also because the function output remains close to zero if the output score vector  $\sigma(\mathbf{o})$  is similar to the label vector  $\mathbf{y}$ . However, the ImageNet Attribute dataset [16] suffers from the problem of class imbalance. As a result, the network is found to learn the class frequencies rather than the classes, when cross-entropy loss is used. To overcome this, we define the weighted cross-entropy loss L as

$$L(\mathbf{y}, \sigma(\mathbf{o})) = -\sum_{k=1}^{N} w_k y_k \log \sigma(o_k), \qquad (1)$$

where k runs over N output classes. The weight vector  $\mathbf{w}$  is calculated from the labels of training samples in a batch. Weight of class k,  $w_k$  is defined as

$$w_k = 1 - \frac{1}{B} \sum_{j=1}^{B} y_{kj} , \qquad (2)$$

where B signifies the number of images in a batch. Eq. (1) ensures higher loss weight for classes having more samples and low weights for classes with fewer samples. With a larger batch size and randomly shuffled data, the dynamically computed loss weight vector reflects the actual class distribution on the entire dataset. Batch size of 32 is used for both training and test images.

The size of the input image used in DAN is the same as in VGG-16, i.e.,  $224 \times 224$ . The convolutional layers and the first fully-connected layer of the

network are initialized from VGG-16, as mentioned before. The weights in the last two fully-connected layers are initialized from a Gaussian distribution with mean 0 and standard deviation of 0.005, and the biases are initialized to a constant value of 0.1. The network is trained using the ImageNet Attribute dataset [16].

The training is done in two phases. First, only the fully-connected layers fc6, fcA and fcB are trained. Once the loss computed on the validation set saturates or starts to increase, the preceding block of convolutional layers (till layer  $conv5_{-1}$ ) are unfrozen and finetuned together with the fully-connected layers. The network is trained for 38 epochs. The training is regularized by weight-decay and dropout regularization in the first two fully-connected layers. The base learning rate is set to 0.001 and is decreased further by factors of 10 when the validation-loss does not improve. All other training parameters are set as follows - momentum to 0.9, weight decay to 0.0005, and dropout rate to 0.5.

# 4 Experiments

In this section, we present the quantitative evaluation of the Deep Attribute Network on the ImageNet Attribute and a-Pascal datasets. We implement our method using the Caffe framework [11].

## 4.1 Datasets, Pre-processing and Metrics

**ImageNet Attribute:** The ImageNet Attribute dataset [16] contains low-level attribute annotations of objects. This dataset is a subset of the ImageNet dataset [15] and consists of 9600 images from 384 synsets, each annotated with 25 attributes belonging to four categories: color (*black, blue, brown, gray, green, orange, pink, red, violet, white, yellow*), shape (*long, round, rectangular, square*), pattern (*spotted, striped*), and texture (*furry, smooth, rough, shiny, metallic, vegetation, wooden, wet*). The labels contain three values -1, +1 and 0, indicating negative, positive and ambiguous samples for which the annotators could not agree on one single attribute. We divide the dataset into training, validation, and test sets with 5760, 1920 and 1920 images, respectively.

**a-Pascal:** The a-Pascal dataset [5] is based on the PASCAL VOC 2008 object recognition dataset [4]. The images contain objects from 20 classes and are annotated with 64 attributes, describing shape, material, and high-level object components. We select three shape attributes (2D boxy, 3D boxy, round) and five material attributes (metal, wood, furry, shiny, vegetation) for our evaluation. The shape attributes 2D boxy and 3D boxy are not present in the training dataset. After manually reviewing the images, we found that both attributes are similar in appearance to the rectangle attribute of ImageNet Attribute dataset. Hence, we merge these two attributes and evaluate them as the rectangle attribute. The labels in this dataset are marked as 0 and 1 for negative and positive samples respectively.

**Pre-processing:** The training and validation images contain object level annotations. Therefore, they are pre-processed by cropping the images to the bounding box of the object with an added margin of 10%. Additionally, the mean RGB value of the training set is subtracted from all images. The label values (-1, 0, +1) in the ImageNet Attribute dataset are mapped to (0, 0.5, 1) for proper loss computation.

**Data augmentation:** The pre-processed images are scaled to size  $256 \times 256$ . At runtime, patches of size  $224 \times 224$  are cropped randomly from the scaled training images, and from the center for the validation images. The training set is further augmented by randomly flipping the images horizontally.

**Evaluation metrics:** The metrics we use for evaluation are the Receiver Operator Characteristic (ROC) curve, the Precision-Recall (P-R) curve, area under the ROC curve (ROC-AUC), and the Average Precision (AP) score for the P-R curve. In a multi-label classification setup, AP for multiple classes is reported as the mean average precision (mAP). The mean can be computed in micro or macro mode [18]. In micro mode, the precision and recall are calculated from the overall true/false positive/negative values across all classes. In macro mode, the average of class-wise precision and recall are computed. The macro-average method gives equal weight to each class, whereas the micro-average method gives equal weight to each class, whereas the micro-average method gives equal weight to each class for a sample. The mAP is calculated using the precision, recall values according to  $mAP = \sum_n (R_n - R_{n-1})P_n$ , where  $R_n$  and  $P_n$  denote the recall and precision at  $n^{th}$  threshold.

## 4.2 Attribute Classification Evaluation

We evaluate the classification performance on the test sets of the ImageNet Attribute and a-Pascal datasets. The ambiguous labels in the ImageNet Attribute dataset are converted to positive labels for computing the precision and recall. The P-R and ROC curves for the test datasets are shown in Fig. 3. DAN achieves an overall mAP score of 0.74 in micro mode and a ROC-AUC score of 0.90 for the ImageNet Attribute dataset. According to the group mAP, texture performs the best, followed by color, shape, and pattern. Images from the pattern classes *striped* and *spotted* vary greatly in appearance, as can be seen in Fig. 4. The fraction of ambiguous samples for these two classes is also high, 68% for *striped*, and 76% for *spotted*. This makes the pattern category difficult for our classifier. For the selected shape and texture attributes from the a-Pascal dataset, DAN performs relatively poor by scoring a micro mode mAP of 0.45 and a ROC-AUC score of 0.75. The AP scores of individual classes are investigated further in Fig. 5.

For the ImageNet Attribute dataset, the AP scores are consistent across all texture classes, resulting in a good mAP score for texture. Among the color classes, the two lowest scores are obtained for *pink* and *violet*. This can be attributed to the lack of sufficient training examples. There are only 78 samples for *pink* and 34 samples for *violet* out of 5760 training samples. Fig. 5 also shows

8



Fig. 3: Performance of the Deep Attribute Network on the ImageNet Attribute and a-Pascal datasets. The figure shows the Precision-recall (left column) and the ROC-AUC curves (right column) for all attributes as well as for each attribute group. Within parentheses in the legends, the micro-mode mAP scores (left column) and the ROC-AUC scores (right column) are reported.



Fig. 4: Sample images from the class *striped* and *spotted* of the ImageNet Attribute dataset showing the large variety in appearance. Some images also contain incorrect labels, for example the box and the bottle in the second row of striped examples and the roller coaster and the bird in the first row of spotted examples.

the performance of the individual classes from the a-Pascal dataset. Though the overall mAP score for this dataset is poor, the class-wise AP scores for the attributes *rectangle*, *furry* and *metal* are at par with the ImageNet Attribute dataset. For the rest of the attributes, the object types differ from the training dataset. For example, the attribute *round* in a-Pascal describes part of objects (wheel of *car*, logo on *train*), whereas in ImageNet Attribute it describes the shape of the objects as a whole. The *vegetation* attribute in a-Pascal dataset



Fig. 5: Class-wise Average-Precision on the ImageNet Attribute and a-Pascal datasets.



Fig. 6: Attribute predictions by DAN on sample images from the ImageNet Attribute dataset. For each image, the attributes with top 3 scores are reported. Correct and incorrect predictions are highlighted in black and red respectively. Blue signifies that the prediction is correct, but is considered wrong due to faulty annotations.

#### 10 Soubarna Banik, Mikko Lauri, Simone Frintrop

describes *pottedplant* in indoor setup, which is very different from the training samples that describe plants in the wild. This affects the classification performance for these attributes.

Fig. 6 shows the attribute classification results for a few sample images from the ImageNet Attribute dataset. The figure shows both correct and incorrect predictions. For example, all the top 3 predicted attributes for the sample images (a,b,c,h,i) are correct. The true label *orange* is incorrectly predicted as *brown* in sample image (k). For sample (l), the attribute *violet* is not predicted. The ImageNet Attribute dataset contains a number of samples, where some object attributes are not annotated. Samples (d,e,f,g,j) highlight such attribute predictions in blue.

Evaluation on cropped test data: The images in the ImageNet Attribute dataset are annotated at object level. Hence any predicted attribute that corresponds to the background is considered false in the evaluation, even if it is correct for the background. To understand the true performance, we also crop the test images to the object bounding boxes with an added margin of 10%. Table 1 reports the micro and macro mode mAP for the un-cropped and the cropped images from the ImageNet Attribute test set. The performance consistently improves for all attribute categories by 2 - 7% in micro mode and 4 - 8% in macro mode compared to using the entire image. The lower macro-mAP scores in both cases suggest that the micro-mAP scores are affected by the class imbalance of the dataset.

**Comparison with state-of-the-art methods:** We compare the performance of the Deep Attribute Network to two state-of-the-art approaches - Russakovsky et al. [16] and Razavian et al. [14]. Russakovsky et al. [16] perform low-level object attribute classification by using handcrafted features and SVM classifiers on the ImageNet Attribute dataset. Razavian et al. [14], on the other hand, perform high-level object attribute classification using features extracted from a ConvNet.

We replicate the approach of Razavian et al. [14] for low-level attribute classification and design an experiment named VGG-16+SVM. We use a pre-trained VGG-16 network as a feature extractor and train four SVMs for the four attribute groups with these features on our training set. Similar to Razavian et

Table 1: Micro and macro mode mAP on the ImageNet Attribute dataset for un-cropped and cropped test images.

	mAP	All	Color	Shape	Pattern	Texture
Un-cropped Un-cropped	micro macro	$\begin{array}{c} 0.74 \\ 0.61 \end{array}$	$0.73 \\ 0.54$	$0.62 \\ 0.56$	$0.52 \\ 0.52$	$0.81 \\ 0.75$
Cropped Cropped	micro macro	$\begin{array}{c} 0.78\\ 0.66 \end{array}$	$\begin{array}{c} 0.77 \\ 0.58 \end{array}$	$\begin{array}{c} 0.68 \\ 0.64 \end{array}$	$0.59 \\ 0.59$	$0.83 \\ 0.80$

Table 2: Comparison of ROC-AUC scores of the proposed approaches with stateof-the-art methods on 20 selected attributes from the ImageNet Attribute test set. Classes *blue*, *violet*, *pink*, *square* and *vegetation* are excluded for a fair comparison with [16].

	Russakovsky et al. [16]	VGG-16+SVM	DAN	DAN-ResNet
Color	0.87	0.81	0.90	0.91
Shape	0.83	0.77	0.87	0.87
Pattern	0.63	0.60	0.81	0.82
Texture	0.77	0.84	0.92	0.91

al.'s set up [14], we use the features from the first fully connected layer of the ConvNet, follow L2 normalization on the features, and use Radial Basis Function (RBF) kernel in the SVMs. The VGG-16+SVM experiment acts as a bridge between the method of Russakovsky et al. involving handcrafted features and the ConvNet based approach proposed in this paper.

We also test our proposed method by using the ResNet-50 ConvNet [7] in place of VGG-16 in order to verify the validity of our approach. We replace the last layer of ResNet with a new fully connected layer containing 25 output nodes and train it for 16 epochs on our training set. All other training parameters are same as DAN. This network is referred to as DAN-ResNet.

Russakovsky et al. [16] exclude the ambiguous records, and also the classes blue, violet, pink, square and vegetation due to insufficient training samples (< 75 positive samples). In order to conduct a fair comparison, we also exclude these classes in the following experiment. As our models jointly learn all attributes, filtering ambiguous records would remove records with even a single ambiguous class label ( $\approx 96\%$  records). Hence, the ambiguous samples are not excluded for DAN and DAN-ResNet. For VGG-16+SVM, the ambiguous labels are converted to positive labels. The ROC-AUC scores of Russakovsky et al.'s method [16], VGG-16+SVM, DAN and DAN-ResNet on the ImageNet Attribute test set are shown in Table 2. The networks based on our proposed approach (DAN and DAN-ResNet) outperform Russakovsky et al.'s approach and VGG-16+SVM by 3 - 18% and 8 - 21% respectively. Particularly, the pattern attribute group shows a major improvement. The results indicate the advantages of learned features over handcrafted features. The relatively poor performance of VGG-16+SVM suggests that features extracted from object classification ConvNet are not sufficient for the attribute classification task and need further fine-tuning.

In a multi-label setup, ROC-AUC score gets biased by the naturally occurring large number of true negatives, as in the case of the ImageNet Attribute dataset. The mAP metric does not suffer from this bias and hence better represents a model's performance. Table 3 shows the micro-mode mAP scores of VGG-16+SVM, DAN and DAN-ResNet on the ImageNet Attribute test set. Russakovsky et al. only report the ROC-AUC score [16]. As their trained models are not available, the mAP scores could not be computed for the same. The

### 12 Soubarna Banik, Mikko Lauri, Simone Frintrop

	VGG-16+SVM		DAN		DAN-ResNet	
#attributes	20	25	20	25	20	25
Color	0.52	0.51	0.75	0.73	0.75	0.74
Pattern	$0.43 \\ 0.18$	0.41	$0.00 \\ 0.52$	$0.02 \\ 0.52$	0.00 0.53	0.03 0.53
Texture	0.64	0.63	0.82	0.81	0.79	0.77

Table 3: Comparison of micro-mode mAP scores of the proposed approaches with state-of-the-art methods on the ImageNet Attribute test set.

mAP scores are calculated for the 20 attributes as in Table 2 and also for all 25 attributes. When the 5 difficult classes (*blue*, *violet*, *pink*, *square* and *vegetation*) are included, the performance decreases consistently. The similar performance of DAN and DAN-ResNet in both Table 2 and 3 demonstrates the validity of our approach.

**Object localization using attributes:** It is interesting to see if objects can be localized by using their attributes. In the past, attention maps generated from object classification ConvNets have been used for locating objects [22]. An attention map shows the highly activated regions of the image which are responsible for a particular classification decision. As part of a preliminary study on attribute localization, we use an algorithm called Excitation Backprop[22] for generating attention maps from DAN. To do so, we provide the attribute to be localized as a feedback signal to the network. Using this approach, attention maps are generated for a sample image for the attributes *red* and *brown*, as shown in Fig. 7. It can be seen that the location of the activation changes depending on the feedback attribute. The red bowl in the input image, though being partially occluded, can be detected using the attribute *red*. This result hints at the promise



Fig. 7: Attention maps generated using Excitation Backprop [22] for attributes *red* and *brown* for a sample input image (shown in the left most column).

of using attributes for localizing objects and we will investigate this approach further in our future work.

# 5 Conclusion

In this paper, we have presented a convolutional neural network for low-level object attribute classification. The novelty of this method is its simple architecture. The network is trained from an object classification ConvNet using transfer learning. It learns 25 attributes belonging to the attribute groups *color*, *shape*, *pattern* and *texture*. Experimental results show that our method outperforms state-of-the-art methods. In the future, we intend to extend this model for attribute localization, which can be helpful in a cluttered or unknown environment, and for unknown objects. The attribute specific attention maps shown in section 4 indicates the plausibility of this future work.

# References

- Abdulnabi, A.H., Wang, G., Lu, J., Jia, K.: Multi-task CNN Model for Attribute Prediction. IEEE Transactions on Multimedia, 2015 17(11), 1949–1959 (2015)
- 2. Chung, J., Lee, D., Seo, Y., Yoo, C.D.: Deep Attribute Networks. In: NIPS Workshop on Deep Learning and Unsupervised Feature Learning, 2012. vol. 3 (2012)
- Cimpoi, M., Maji, S., Vedaldi, A.: Deep Filter Banks for Texture Recognition and Segmentation. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 2015. pp. 3828–3836. IEEE (2015)
- Everingham, M., Van Gool, L., Williams, C.K.I., Winn, J., Zisserman, A.: The PASCAL Visual Object Classes Challenge 2008 (VOC2008) Results. http://www.pascal-network.org/challenges/VOC/voc2008/workshop/index.html
- Farhadi, A., Endres, I., Hoiem, D., Forsyth, D.: Describing Objects by their Attributes. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 2009. pp. 1778–1785. IEEE (2009)
- Ferrari, V., Zisserman, A.: Learning Visual Attributes. In: Advances in Neural Information Processing Systems (NIPS), 2008. pp. 433–440 (2008)
- He, K., Zhang, X., Ren, S., Sun, J.: Deep Residual Learning for Image Recognition. In: Proceedings of the IEEE conference on Computer Vision and Pattern Recognition (CVPR), 2016. pp. 770–778 (2016)
- Hu, C., Bai, X., Qi, L., Chen, P., Xue, G., Mei, L.: Vehicle Color Recognition with Spatial Pyramid Deep Learning. Proceedings of the IEEE Transactions on Intelligent Transportation Systems 16(5), 2925–2934 (2015)
- Huang, S., Elhoseiny, M., Elgammal, A., Yang, D.: Learning Hypergraphregularized Attribute Predictors. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 2015. pp. 409–417. IEEE (2015)
- Jayaraman, D., Sha, F., Grauman, K.: Decorrelating Semantic Visual Attributes by resisting the urge to share. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 2014. pp. 1629–1636 (2014)
- Jia, Y., Shelhamer, E., Donahue, J., Karayev, S., Long, J., Girshick, R., Guadarrama, S., Darrell, T.: Caffe: Convolutional Architecture for Fast Feature Embedding. arXiv preprint arXiv:1408.5093 (2014)

- 14 Soubarna Banik, Mikko Lauri, Simone Frintrop
- Lampert, C.H., Nickisch, H., Harmeling, S.: Learning to Detect Unseen Object Classes by Between-class Attribute Transfer. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 2009. pp. 951–958. IEEE (2009)
- Liu, Z., Luo, P., Wang, X., Tang, X.: Deep Learning Face Attributes in the Wild. In: Proceedings of the IEEE International Conference on Computer Vision (ICCV), 2015. pp. 3730–3738 (2015)
- Razavian, A.S., Azizpour, H., Sullivan, J., Carlsson, S.: CNN features off-the-shelf: An astounding baseline for recognition. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 2014. pp. 512–519. IEEE (2014)
- Russakovsky, O., Deng, J., Su, H., Krause, J., Satheesh, S., Ma, S., Huang, Z., Karpathy, A., Khosla, A., Bernstein, M., et al.: Imagenet Large scale Visual Recognition Challenge. International Journal of Computer Vision 115(3), 211–252 (2015)
- Russakovsky, O., Fei-Fei, L.: Attribute Learning in Large-scale Datasets. In: Proceedings of the European Conference on Computer Vision (ECCV), 2010. pp. 1–14. Springer (2010)
- Simonyan, K., Zisserman, A.: Very Deep Convolutional Networks for Large-scale Image Recognition. arXiv preprint arXiv:1409.1556 (2014)
- Tsoumakas, G., Katakis, I., Vlahavas, I.: Mining Multi-label Data. In: Data mining and Knowledge Discovery Handbook, pp. 667–685. Springer (2009)
- Wang, J., Yang, Y., Mao, J., Huang, Z., Huang, C., Xu, W.: CNN-RNN: A Unified Framework for Multi-label Image Classification. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 2016. pp. 2285– 2294. IEEE (2016)
- Wu, F., Wang, Z., Lu, W., Li, X., Yang, Y., Luo, J., Zhuang, Y.: Regularized Deep Belief Network for Image Attribute Detection. Proceedings of the IEEE Transactions on Circuits and Systems for Video Technology, 2017 27(7), 1464–1477 (2017)
- Yosinski, J., Clune, J., Bengio, Y., Lipson, H.: How transferable are features in Deep Neural Networks? In: Advances in Neural Information Processing Systems. pp. 3320–3328 (2014)
- Zhang, J., Lin, Z., Brandt, J., Shen, X., Sclaroff, S.: Top-down Neural Attention by Excitation Backprop. In: Proceedings of the European Conference on Computer Vision (ECCV), 2016. pp. 543–559. Springer (2016)
- Zhang, N., Paluri, M., Ranzato, M., Darrell, T., Bourdev, L.: Panda: Pose aligned Networks for Deep Attribute Modeling. In: Proceedings of the IEEE conference on Computer Vision and Pattern Recognition (CVPR), 2014. pp. 1637–1644 (2014)