

Localizing by Describing: Attribute-Guided Attention Localization for Fine-Grained Recognition

Xiao Liu,* Jiang Wang,* Shilei Wen, Errui Ding, Yuanqing Lin

Baidu Research

{liuxiao12, wenshilei, dingerrui, linyuanqing}@baidu.com

wangjiangb@gmail.com

Abstract

A key challenge in fine-grained recognition is how to find and represent discriminative local regions. Recent attention models are capable of learning discriminative region localizers only from category labels with reinforcement learning. However, not utilizing any explicit part information, they are not able to accurately find multiple distinctive regions. In this work, we introduce an attribute-guided attention localization scheme where the local region localizers are learned under the guidance of part attribute descriptions. By designing a novel reward strategy, we are able to learn to locate regions that are spatially and semantically distinctive with reinforcement learning algorithm. The attribute labeling requirement of the scheme is more amenable than the accurate part location annotation required by traditional part-based fine-grained recognition methods. Experimental results on the CUB-200-2011 dataset (Wah et al. 2011) demonstrate the superiority of the proposed scheme on both fine-grained recognition and attribute recognition.

Introduction

Humans heavily rely on subtle local visual cues to distinguish fine-grained object categories. For example in Figure (a), human experts differentiate a summer tanager and a scarlet tanager by the color of wing and tail. In order to build human-level fine-grained recognition AI systems, it is also essential to locate discriminative object parts and learn local visual representation for these parts.

State-of-the-art fine-grained recognition methods (Zhang et al. 2014; Liu, Shen, and Hengel 2015; Wang et al. 2014; Krause et al. 2015b; 2015a; Gao et al. 2015; Zhang et al. 2016a) either rely on manually labeled parts to train part detectors in a fully supervised manner, or employ reinforcement learning or spatial-transformer-based attention models (Sermanet, Frome, and Real 2015; Liu et al. 2016b; Zhao et al. 2016) to locate object parts with object category annotations in a weakly supervised manner. However, both types of methods have major practical limitations. Fully supervised methods need time-consuming, error-probing manual object part labeling process, while object labels alone as a supervision signal is generally too weak to reliably

locate multiple discriminative object parts. For example in Figure (b), the method (Liu et al. 2016b) fails to locate the tail as attention regions.

Humans have a remarkable ability to learn to locate object parts from multiple sources of information. Aside from strongly labeled object part location and weakly labeled object categories, part descriptions, such as “red wing”, also plays an important part in the development of object parts locating ability. In their early ages, children learn to recognize parts and locate object parts by reading or listening to part descriptions. Part descriptions do not require time-consuming manually part location labeling, and it is much stronger than object category labels. We call it *part attribute*.

Inspired by this capability, we propose a part attribute-guided attention localization scheme for fine-grained recognition. Using part attributes as a weak supervision training signal, reinforcement learning is able to learn part-specific optimal localization strategies given the same image as environment state. Based on this intuition, it is reasonable to expect that distinctive part localizers could be learned as strategies of looking for and describing the appearances of different parts. In the proposed scheme, multiple fully convolutional attention localization networks are trained. Each network predicts the attribute values of a part. We design a novel reward strategy for learning part localizers and part attribute predictors.

Part attribute-guided attention localization networks can more accurately locate object parts (Figure (c)). More importantly, using the part locations and appearance features from part-attribute guided attention localization networks leads to considerable performance improvement on fine-grained recognition, as demonstrated on the CUB-200-2011 dataset (Wah et al. 2011). Moreover, part attribute can be acquired in large scale via either human labeling or data mining techniques. It has been successfully employed for image recognition (Parikh and Grauman 2011; Akata et al. 2013; Hwang and Sigal 2014), image retrieval (Huang et al. 2015), face localization (Liu et al. 2015) and image generation (Yan et al. 2015).

Attention model (Mnih et al. 2014; Ba, Mnih, and Kavukcuoglu 2015) has been widely used in image caption (Xu and Saenko 2015; Xu et al. 2015), image recognition (Zhang et al. 2016b; Seo et al. 2016), action recognition (Wang et al. 2016), person re-identification (Liu et

*Equal contribution.

al. 2016a) and image parsing (Fan et al. 2016). To our best knowledge, no one has utilized part attribute to guide the learning of visual attention in fine-grained recognition.

Attribute-Guided Attention Localization

The architecture of the proposed scheme is shown in Figure 2. Attribute descriptions are only used for learning part localizers during training. They are not used in the testing stage.

In the training stage, a fully-convolutional attention localization network (Liu et al. 2016b) is learned for each part. In contrast with (Liu et al. 2016b), the task of the fully-convolutional attention localization network is to learn where to look for better part attribute prediction rather than predicting the category of the entire object.

For testing, we extract features from both part regions located by the network and the entire image and concatenate them as a joint representation. The joint representation is utilized to predict the image category.

Problem Formulation

Given N training images $\{x_1, x_2, \dots, x_N\}$, and their object labels $\{l_1, l_2, \dots, l_N\}$. Our goal is to train a model for classifying each image x_i as its ground-truth label l_i .

Fine-grained objects have P parts. *Localizing by describing* problem finds a policy to locate the parts and to classify images based on these part locations and their features. The objective can be formulated as:

$$\mathcal{L}_0 = \sum_i L(\mathcal{G}(x_i, \mathcal{F}_1(x_i), \dots, \mathcal{F}_P(x_i)), l_i), \quad (1)$$

where $\mathcal{F}_p(\cdot)$ is a function that finds the location of part p and crops its image region. The crop size of each part is manually defined. $\mathcal{G}(\cdot)$ is a deep convolutional neural network classifier that outputs the probability of each category given the whole image and the cropped image regions for all the parts. $L(\cdot)$ is cross-entropy loss function measuring the quality of the classification.

Precisely predicting part locations using only the image labels is very challenging. In *localizing by describing*, we localize object parts with the help of visual attribute descriptions. The visual attribute is a set of binary annotations whose value correlates with an aspect of the object appearance. A local attribute is generally related to the appearance of an object part. The attribute description of the p -th part is denoted as $\{Y_1^p, Y_2^p, \dots, Y_N^p\}$. Each part description Y_i^p is a binary vector: $[y_{i,1}^p, y_{i,2}^p, \dots, y_{i,K_p}^p]$, where each element $y_{i,k}^p$ indicates whether an attribute exists in the i -th image, and K_p is the number of attributes for the p -th part. We aim to learn better part localizers \mathcal{F}_p using part attribute annotation. An auxiliary part attribute classification loss is proposed to facilitate the learning of part localizers

$$\mathcal{L}_p = \sum_i L'(\mathcal{T}_p(\mathcal{F}_p(x_i)), Y_i^p), \quad (2)$$

where $\mathcal{T}_p(\cdot) = [\mathcal{T}_{p,1}(\cdot), \dots, \mathcal{T}_{p,K_p}(\cdot)]$ is a multi-label attribute prediction function of the p -th part, and each $\mathcal{T}_{p,k}(\cdot)$

indicates the predicted probability of the k -th attribute of part p . $L'(\cdot)$ is a multi-label cross-entropy loss

$$L'(X_i, Y_i^p) = - \sum_k [y_{i,k}^p \log \mathcal{T}_{p,k}(\mathcal{F}_p(x_i)) + (1 - y_{i,k}^p) \log(1 - \mathcal{T}_{p,k}(\mathcal{F}_p(x_i)))]. \quad (3)$$

The assumption is that the part localizers that help predict part attributes are also beneficial to the prediction of the whole object. We use Eq. (3) as an auxiliary loss to learn localizer $\mathcal{F}_p(\cdot)$ that optimizes Eq.(1).

Training of Localizers

Given images and attribute descriptions, we jointly learn a part localizer and a multi-label predictor for each part such that the predictor uses the selected local region for attribute prediction.

Since the localization operation is non-differential, we employ reinforcement learning algorithm (Williams 1992) to learn the part localizers and multi-label predictors. For reinforcement learning algorithm, the policy function is the part localizer function $\mathcal{F}_p(\cdot)$; the state is the cropped local image patch $\mathcal{F}_p(x_i)$ and the whole image x_i ; the reward function measures the quality of the part attribute.

The objective function of the reinforcement learning algorithm for part p is

$$J_p(\mathcal{F}_p, \mathcal{T}_p) = \sum_i [\mathbb{E}_{\mathcal{F}_p}(R_{p,i}) - \lambda L'(\mathcal{T}_p(\mathcal{F}_p(x_i)), Y_i^p)], \quad (4)$$

where the reward

$$\mathbb{E}_{\mathcal{F}_p}(R_{p,i}) = \int_{s_{p,i}} p_{\mathcal{F}_p}(s_{p,i}|x_i) r(s_{p,i}) d(s_{p,i}) \quad (5)$$

is the expected reward of the selected region for the p -th part of the i -th image. $s_{p,i}$ indicates a selected region, $p_{\mathcal{F}_p}(s_{p,i}|x_i)$ is the probability that the localizer select the region $s_{p,i}$. $r(s_{p,i})$ is a reward function to evaluate the contribution of the selected region $s_{p,i}$ to attribute prediction.

Previous methods (Sermanet, Frome, and Real 2015; Liu et al. 2016b) choose the reward function $r(s_{p,i})$ to be 1 only when the image is correctly classified. However, since our algorithm predicts multiple attribute values for a part, it is too strict to enforce all the attributes are correctly predicted. Therefore, we consider an alternative reward strategy. A selected region has reward 1 if both of the following criteria are satisfied: 1) it achieves lower attribute classification loss than most other regions in the same image, *i.e.*, its prediction loss ranks top- η lowest among the M sampled regions of the image. 2) it achieves lower attribute classification loss than most other regions in the same mini-batch, *i.e.*, its prediction loss is lower than half of the average loss of all the regions in the mini-batch.

Following (Liu et al. 2016b), we learn fully convolutional attention localization networks as part localizers. Since both parts of the objective function Eq. (4) are differentiable, REINFORCE algorithm (Williams 1992) is applied to compute

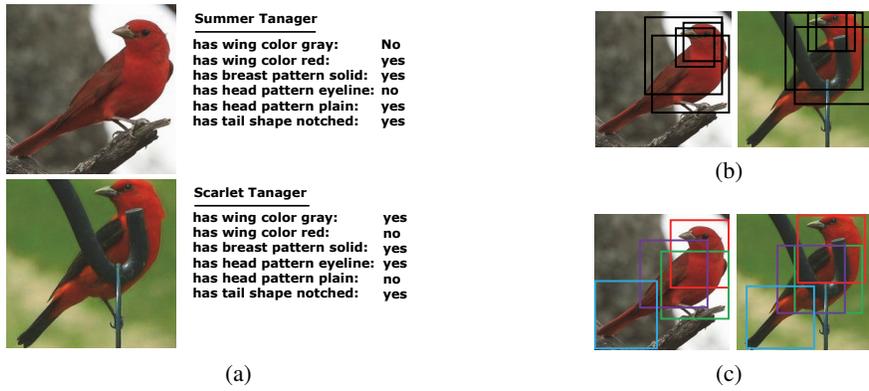


Figure 1: (a) Two types of birds with part-attribute descriptions. Human experts differentiate a summer tanager and a scarlet tanager by the color of wing and tail. (b) (Liu et al. 2016b) fails to locate the tail part. (c) Part-attribute-guided attention localization networks can locate object parts more accurately. The red, green, purple and blue bounding boxes localize head, breast, wing and tail, respectively (best viewed in color).

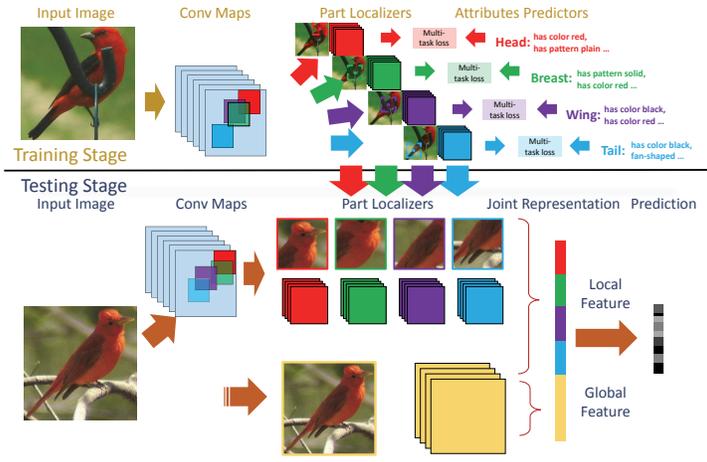


Figure 2: An overview architecture of the proposed scheme. The correspondence of color and part is the same as Fig. 1. The upper part shows training stage, and the lower part shows the testing stage. In the training stage, multiple part localizers are trained under the guidance of attribute descriptions. In the testing stage, features extracted from the selected part regions and the entire image are combined into a joint representation for category prediction.

the policy gradient to optimize the objective function:

$$\begin{aligned}
 & \partial_{\mathcal{F}_p} \mathbb{E}_{\mathcal{F}_p}(R_{p,i}) \\
 = & \int_{s_{p,i}} p_{\mathcal{F}_p}(s_{p,i}|x_i) \partial_{\mathcal{F}_p} \log[p_{\mathcal{F}_p}(s_{p,i}|x_j)r(s_{p,i})] d(s_{p,i}) \\
 \approx & \frac{1}{M} \sum_{m=1}^M \partial_{\mathcal{F}_p} \log[p_{\mathcal{F}_p}(s_{p,i}^m|x_i)r(s_{p,i}^m)]
 \end{aligned} \tag{6}$$

where $s_{i,j}^m \sim p_{\mathcal{F}_i}(\cdot|x_j)$ is the local image regions sampled according to localizer policy $p_{\mathcal{F}_i}(s_{i,j}^m|x_j)$. M local regions of the same image are sampled in a mini-batch. We list the learning algorithm in Algorithm 1.

Training of Classifiers

After the local region localizers are trained, we re-trained the attribute prediction models using up-scaled local regions. When re-trained, the attribute predictors takes up-scaled local regions from the part localizers to predict the attributes.

To combine global and local information, we extract features from all the part regions and the entire image and concatenate them to form a joint representation. The joint representation is used to predict image category. In details, we first train classifiers with each individual part region to capture the appearance details of local parts for fine-grained recognition. A classifier utilizing the entire image is also trained for global information. We then concatenate features extracted from all the parts and the entire image as a joint representation, and we use a linear layer to combine the fea-

Algorithm 1 Localizing by describing algorithm:

Input: training images $\{x_1, x_2, \dots, x_N\}$, attribute descriptions of each part $\{Y_1^p, Y_2^p, \dots, Y_N^p\}$.
Output: part localization function \mathcal{F} , multi-label attribute prediction function \mathcal{T} .

- 1: **for** each part p **do**
- 2: Initialize \mathcal{F}_p and \mathcal{T}_p .
- 3: **repeat**
- 4: Randomly sample H images.
- 5: **for** each image x_i **do**
- 6: Sample $s_{p,i}^m \sim p_{\mathcal{F}_p}(\cdot|x_i)$
- 7: **end for**
- 8: **for** each local region $s_{p,i}^m$ **do**
- 9: Calculate $l_{p,i}^m = L'(\mathcal{T}_p(s_{p,i}^m), Y_i^p)$.
- 10: **end for**
- 11: Calculate $\hat{L} = \sum l_{p,i}^m$.
- 12: **for** each image x_i **do**
- 13: Sort $l_{p,i}^m$ in ascending order.
- 14: **for** each local region $s_{p,i}^m$ **do**
- 15: **if** $l_{p,i}^m$ is in the top- η position of the sorted list,
 and $l_{p,i}^m < 0.5\hat{L}$ **then**
- 16: Set reward $r(s_{p,i}^m) = 1$.
- 17: **else**
- 18: Set reward $r(s_{p,i}^m) = 0$.
- 19: **end if**
- 20: **end for**
- 21: **end for**
- 22: Calculate the gradient of \mathcal{F}_p according to (6).
- 23: Calculate the gradient of \mathcal{T}_p .
- 24: Update the parameters of \mathcal{F}_p and \mathcal{T}_p .
- 25: **until** converge
- 26: **end for**

tures.

Prediction

The prediction process is illustrated in the lower part of Fig. 2. We localize and crop each part using its localizer $\mathcal{F}_p(x_i)$, and each part region is resized to high resolution for feature extraction. Features from all part regions as well as the entire image are concatenated as a joint representation. A linear classification layer is utilized to make the final category prediction.

We also attempt to use attribute prediction results to help fine-grained object recognition, or model the geometric relationship of the parts using recurrent convolutional operation. However, we find neither of the approaches achieve notable improvements in our experiments. Detailed experimental results and setup can be found in the experimental section.

Experiments

We conduct experiments on the CUB-200-2011 datasets (Wah et al. 2011). The dataset contains 11,788 images of 200 bird categories, where 5,994 images are for training, and the rest 5,794 images are for testing. In addition to the category label, 15 part locations, 312 binary attributes and a

tight bounding box of the bird is provided for each image. Examples of images and attributes in this dataset are shown in Figure (a).

Implementation Details

We evaluate the proposed scheme in two scenarios: “with BB” where the object bounding box is utilized during training and testing, and “without BB” where the object bounding-box is not utilized.

We choose 4 parts, *i.e.* “head”, “wing”, “breast”, and “tail”, to train local region localizers. The cropping size of these parts are half of the original image. Among all the 312 attributes, if a part name appears in an attribute, then the attribute is considered describing this part. The number of attributes describing the four parts are 29, 24, 19, and 40, respectively.

We utilize ResNet-50 (He et al. 2016) as the visual representation for part localization and feature extraction. In the training stage, we utilize the ROI-pooled feature maps (in the “with BB” setting, crop the image region within the bounding box) to learn multi-label attribute predictions for each part. The output of the “res5c” layer of ResNet-50 is employed as the input of the fully convolutional attention localization networks. The attribute predictors use ROI-pooled feature maps (Girshick 2015) of the fully convolutional attention localizers.

We train the models using Stochastic Gradient Descent (SGD) with momentum of 0.9, epoch number of 150, weight decay of 0.001, and a mini-batch size of 28 on four K40 GPUs. One epoch means all training samples are passed through once. An additional dropout layer with an ratio of 0.5 is added after “res5c”, and the size of “fc15” is changed from 1000 to 200.

The parameters before “res5c” are initialized by the model (He et al. 2016) pretrained on the ImageNet dataset (Deng et al. 2009), and parameters of fc15 are randomly initialized. The initial learning rate is set at 0.0001 and reduced twice with a ratio of 0.1 after 50 and 100 epoches. The learning rate of the last layer (“fc15”) is 10 times larger than other layers.

Our data augmentation is similar to (Szegedy et al. 2014), but we have more types of data augmentation. A training image is first rotated with a random angle between -30° and 30° . A cropping is then applied on the rotated image. The size of the cropped patch is chosen randomly between 25% and 100% of the whole image, and its aspect ratio is chosen randomly between 3/4 and 4/3. AlexNet-style color augmentation (Krizhevsky, Sutskever, and Hinton 2012) is also applied followed by random flip. We finally resize the transformed cropped patch to a 448×448 image as the input of the convolutional neural network.

Part Localization Results

We report our part localization results in Table 1. Percent Correct Parts (PCP) is used as the evaluation metric. A part is determined to be correctly localized if the difference of its predicted location and the ground-truth location is within 1.5 times the ground-truth annotations standard deviation.

Method	head	breast	wing	tail	ave
(Shih et al. 2015)	67.6	77.8	81.3	59.2	71.4
(Liu and Belhumeur 2013)	58.5	67.0	71.6	40.2	59.3
(Liu, Li, and Belhumeur 2014)	72.0	70.5	74.4	46.2	65.8
Ours (without BB)	60.6	79.5	77.5	47.1	66.2
Ours (with BB)	69.3	81.5	80.3	62.5	73.4

Table 1: Part localization results (measured by PCP) on the CUB-200-2011 dataset.

We compare with previous state-of-the-art part localization methods (Shih et al. 2015; Liu and Belhumeur 2013; Liu, Li, and Belhumeur 2014). The strongly supervised method (Shih et al. 2015), which utilizes the precise location of the parts during training, achieves the highest average PCP (71.4) in the “without BB” scenario. Our scheme that does not use any part or object locations achieves the second highest average PCP (66.2) and performs the best for localizing “breast”. The parts are localized much more precisely (66.2 \rightarrow 73.4) when ground-truth bird bounding boxes is used.

Figure 3 provides visualizations of our part localization results. Ground-truth part locations are shown as hollow circles, predicted part locations are shown as solid circles and the selected part regions are shown as thumbnails.

It should be noted that different from (Shih et al. 2015), our scheme does not directly minimize the part localization error but learns part localizers for attribute prediction. Thus, a selected region that is far from manually annotation might better predict some part attributes. For example, our predicted “head” positions are usually in the center of the bird head such that the cropped local region does not lose any information, while the manually annotated “head” positions normally appear on the “forehead”.

Attribute Prediction Results

The attribute prediction results measured by average Area Under the Curve of ROC (AUC) are reported in Table 2. We use AUC instead of accuracy because the data of attributes are highly imbalanced: most attributes are only activated in a very small number of images, but some attributes appear very frequently. For each part, we calculate the AUC of all its attributes, and report the average AUC of them as the AUC of the part.

Directly utilizing the full image as input achieves 76.8, 78.7, 73.1, and 70.4 average AUC for parts “head”, “breast”, “wing” and “tail”, respectively. The overall average AUC of the four parts is 74.3. Using the localized local regions results in slight performance drop (74.3 \rightarrow 73.7) on overall average AUC. The prediction using both the full image and the local attention regions improves the overall average AUC result to 76.6. Bounding boxes of birds are not used for attribute prediction.

Fine-grained Recognition Results

For “without BB” scenario, the baseline ResNet-50 using the whole image achieves 81.7% recognition accuracy. Adding features of two parts (“head” and “breast”) improves the result to 84.7% and combing features of four parts improves

Method	head	breast	wing	tail	total
Full image	76.8	78.7	73.1	70.4	74.3
Attention	77.8	78.3	71.9	68.8	73.7
Image + Attention	80.3	80.7	75.5	72.1	76.6

Table 2: Attribute prediction results (measured by average AUC) on the CUB-200-2011 dataset.

the result to 85.1%. For “with BB” scenario, the baseline achieves 82.3% accuracy. Combing features of two parts improves the result to 84.9%, and combing all the features improves the accuracy to 85.3%. Using features of GoogLeNet (Ioffe and Szegedy 2015) and VGGNet (Simonyan and Zisserman 2014) achieves 85.1% and 84.0%, respectively.

We carry out further experiments to explore using the attributes for better recognition. In the “full image + attribute value” experiment, we concatenate 112 binary attribute labels with the original visual feature of the whole image to predict the bird category. In the “full image + attribute feature” experiment, we concatenate the visual features of the 5 part attribute prediction models: one is the original full image model, and the other four models are fine-tuned for predicting the attributes of the four parts.

As Table 3 shows, directly combining attribute values does not improve recognition accuracy compared with the baseline. Combing attribute features leads to marginal improvements (81.7% \rightarrow 82.5% and 82.3% \rightarrow 82.9%), because we find the attribute predictions are usually noisy due to the inherent difficulty of predicting some local attributes. By combining features from the full image model, four part-based models and the attribute prediction models, we achieve the 85.4% and 85.5% for the “without BB” and “with BB” scenarios, respectively. We also explore jointly training the part localizers using a recurrent convolutional neural network model as geometric regularization of part locations (Tompson et al. 2014), but we find the accuracy improvement is negligible (85.1% \rightarrow 85.2% without BB). After examining the data, we find birds poses are too diverse to learn an effective geometric model from limited amount of data.

We compare with previous state-of-the-art methods on this dataset and summarize the recognition results in Table 3. (Zhang et al. 2014) train part-based R-CNN model to detect the head and body of the bird. The method relies on part location annotation during training. Our scheme outperforms (Zhang et al. 2014) and other part-based models (Simon and Rodner 2015; Zhang et al. 2015) by a large margin without requiring strongly supervised part location annota-



Figure 3: Visualizations of part localization results. Ground-truth part locations are shown as hollow circles; predicted part locations are shown as solid circles; selected part regions are shown as thumbnails. Different parts are shown in different colors. No ground-truth bird bounding boxes or part locations are used during training or testing time in the proposed method.

Method	Acc without BB(%)	Acc with BB(%)
(Lin, RoyChowdhury, and Maji 2015)	84.1	85.1
(Krause et al. 2015a)	82.0	82.8
(Zhang et al. 2014)	73.9	76.4
(Liu et al. 2016b)	82.0	84.3
(Jaderberg et al. 2015)	84.1	-
(Zhang et al. 2015)	75.0	-
(Simon and Rodner 2015)	81.0	-
Full Image	81.7	82.3
Full Image + 2×parts	84.7	84.9
Full Image + 4×parts	85.1	85.3
Full Image + attribute value	81.7	82.2
Full Image + attribute feature	82.5	82.9
Full Image + 4×parts + attribute feature	85.4	85.5

Table 3: Recognition results on the CUB-200-2011 dataset with different settings.

tion. (Lin, RoyChowdhury, and Maji 2015) construct high dimensional bilinear feature vectors, and achieve 84.1% and 85.1% accuracy for the “without BB” and “with BB” scenarios, respectively. (Krause et al. 2015a) learn and combine multiple latent parts in a weakly supervised manner, and achieve 82.0% and 82.8% accuracy for the “without BB” and “with BB” scenarios, respectively. (Liu et al. 2016b) utilize the fully convolutional attention localization networks to select two local parts for model combination. The accuracy is 82.0% without bounding box and 84.3% with bounding box, while our accuracy is 84.7% without bounding box and 84.9% with bounding box by combing features from two local parts (“head” and “breast”). Localizing distinctive parts leads to better recognition accuracy. Similarly, (Jaderberg et al. 2015) combine features of four local parts and achieve an accuracy of 84.1% without using bounding box. Our scheme using the same number of parts outperforms it by 1% (84.1% → 85.1%).

Reward Strategy Visualization

The rewards during reinforcement learning algorithm are visualized in Figure 4. From left to right, we show the heatmaps of 1st, 40-th, 80-th, 120-th, 160-th, and 200-th iterations for multi-label attribute prediction loss, rewards and the output probability of the localizer. As can be seen, after an initial divergence on localizer probability map, the out-

put of the localizer converges to the “head” position during training as expected.

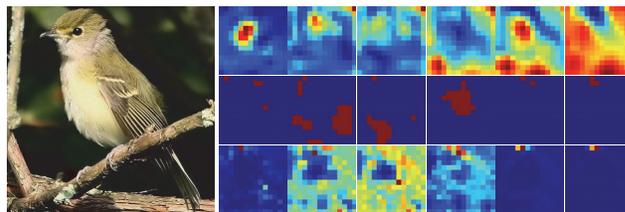


Figure 4: Our algorithm uses 200 iterations to locate “head” of the bird in the left part. The top row in the right part shows the multi-label attribute prediction loss at different positions, and a lighter position has higher loss. The middle row shows the rewards at different positions. The localizer is encouraged to focus on the light position. The bottom row shows the probability map of the localizer. Lighter positions indicate larger probability of localization.

Conclusion

In this paper, we present an attention localization scheme for fine-grained recognition that learns part localizers from its attribute descriptions. An efficient reinforcement learn-

ing scheme is proposed for this task. The proposed scheme consists of a reward function for encouraging different part localizers to capture complementary information. It is also highly computationally efficient when the number of attributes and parts is large. Comprehensive experiments show that our scheme obtains good part localization, improves attribute prediction, and demonstrates significant performance improvement on fine-grained recognition. In the future, we will continue our efforts to improve the models of geometric part location regularization.

References

- Akata, Z.; Perronnin, F.; Harchaoui, Z.; and Schmid, C. 2013. Label-embedding for attribute-based classification. *Proc. ECCV*.
- Ba, J.; Mnih, V.; and Kavukcuoglu, K. 2015. Multiple object recognition with visual attention. *arXiv:1412.7755v2*.
- Deng, J.; Dong, W.; Socher, R.; Li, L.; Li, K.; and Fei-Fei, L. 2009. Imagenet: A large-scale hierarchical image database. *Proc. CVPR*.
- Fan, H.; Mei, X.; Prokhorov, D.; and Ling, H. 2016. Hierarchical attention network for action recognition in videos. *arXiv:1607.02537*.
- Gao, Y.; Beijbom, O.; Zhang, N.; and Darrell, T. 2015. Compact bilinear pooling. *arXiv:1511.06062*.
- Girshick, R. 2015. Fast r-cnn. *arXiv preprint arXiv:1504.08083*.
- He, K.; Zhang, X.; Ren, S.; and Sun, J. 2016. Deep residual learning for image recognition. *Proc. CVPR*.
- Huang, J.; Feris, R. S.; Chen, Q.; and Yan, S. 2015. Cross-domain image retrieval with a dual attribute-aware ranking network. *Proc. ICCV*.
- Hwang, S. J., and Sigal, L. 2014. A unified semantic embedding: Relating taxonomies and attributes. *Proc. NIPS*.
- Ioffe, S., and Szegedy, C. 2015. Batch normalization: accelerating deep network training by reducing internal covariate shift. *Proc. ICML*.
- Jaderberg, M.; Simonyan, K.; Zisserman, A.; and Kavukcuoglu, K. 2015. Spatial transformer networks. *Proc. NIPS*.
- Krause, J.; Jin, H.; Yang, J.; and Fei-Fei, L. 2015a. Fine-grained recognition without part annotations. *Proc. CVPR*.
- Krause, J.; Sapp, B.; Howard, A.; Zhou, H.; Toshev, A.; Duerig, T.; Philbin, J.; and Fei-Fei, L. 2015b. The unreasonable effectiveness of noisy data for fine-grained recognition. *arXiv:1511.06789v1*.
- Krizhevsky, A.; Sutskever, I.; and Hinton, G. E. 2012. ImageNet classification with deep convolutional neural networks. *Proc. NIPS*.
- Lin, T.; RoyChowdhury, A.; and Maji, S. 2015. Bilinear cnn models for fine-grained visual recognition. *Proc. ICCV*.
- Liu, J., and Belhumeur, P. 2013. Bird part localization using exemplar-based models with enforced pose and subcategory consistency. *Proc. ICCV*.
- Liu, Z.; Luo, P.; Wang, X.; and Tang, X. 2015. Deep learning face attributes in the wild. *Proc. ICCV*.
- Liu, H.; Feng, J.; Qi, M.; Jiang, J.; and Yan, S. 2016a. End-to-end comparative attention networks for person re-identification. *1606.04404*.
- Liu, X.; Xia, T.; Wang, J.; and Lin, Y. 2016b. Fully convolutional attention localization networks: efficient attention localization for fine-grained recognition. *arXiv: 1603.06765*.
- Liu, J.; Li, Y.; and Belhumeur, P. 2014. Part-pair representation for part localization. *Proc. ECCV*.
- Liu, L.; Shen, C.; and Hengel, A. 2015. The treasure beneath convolutional layers: cross-convolutional-layer pooling for image classification. *Proc. CVPR*.
- Mnih, V.; Heess, N.; Graves, A.; and Kavukcuoglu, K. 2014. Recurrent models of visual attention. *Proc. NIPS*.
- Parikh, D., and Grauman, K. 2011. Relative attributes. *Proc. ICCV*.
- Seo, P. H.; Lin, Z.; Cohen, S.; Shen, X.; and Han, B. 2016. Hierarchical attention networks. *arXiv:1606.02393*.
- Sermanet, P.; Frome, A.; and Real, E. 2015. Attention for fine-grained categorization. *arXiv:1412.7054v3*.
- Shih, K.; Mallya, A.; Singh, S.; and Hoiem, D. 2015. Part-localization using multi-proposal consensus for fine-grained categorization. *Proc. BMVC*.
- Simon, M., and Rodner, E. 2015. Neural activation constellations: Unsupervised part model discovery with convolutional networks. *ICCV*.
- Simonyan, K., and Zisserman, A. 2014. Very deep convolutional networks for large-scale image recognition. *arXiv:1409.1556*.
- Szegedy, C.; Liu, W.; Jia, Y.; Sermanet, P.; Reed, S.; Anguelov, D.; Erhan, D.; Vanhoucke, V.; and Rabinovich, A. 2014. Going deeper with convolutions. *arXiv:1409.4842*.
- Tompson, J.; Jain, A.; LeCun, Y.; and Bregler, C. 2014. Joint training of a convolutional network and a graphical model for human pose estimation. *Proc. NIPS*.
- Wah, C.; Branson, S.; Welinder, P.; Perona, P.; and Belongie, S. 2011. The caltech-ucsd birds-200-2011 dataset.
- Wang, X.; Yang, T.; Chen, G.; and Lin, Y. 2014. Object-centric sampling for fine-grained image classification. *arXiv:1412.3161*.
- Wang, Y.; Wang, S.; Tang, J.; O'Hare, N.; Chang, Y.; and Li, B. 2016. Hierarchical attention network for action recognition in videos. *arXiv:1607.06416*.
- Williams, R. 1992. Simple statistical gradient-following algorithms for connectionist reinforcement learning. *Machine Learning*.
- Xu, H., and Saenko, K. 2015. Ask, attend and answer: Exploring question-guided spatial attention for visual question answering. *arXiv:1511.05234*.
- Xu, K.; Ba, J.; Kiros, R.; Cho, K.; Courville, A.; Salakhutdinov, R.; Zemel, R.; and Bengio, Y. 2015. Show, attend and tell: Neural image caption generation with visual attention. *arXiv:1502.03044*.
- Yan, X.; Yang, J.; Sohn, K.; and Lee, H. 2015. Attribute2image: Conditional image generation from visual attributes. *arXiv:1512.00570*.
- Zhang, N.; Donahue, J.; Girshick, R.; and Darrell, T. 2014. Part-based r-cnns for fine-grained category detection. *Proc. ECCV*.
- Zhang, Y.; Shen, X.; Wu, J.; Cai, J.; Lu, J.; Nguyen, V.-A.; and Do, M. N. 2015. Weakly supervised fine-grained image categorization. *arXiv:1504.04943*.
- Zhang, H.; Xu, T.; Elhoseiny, M.; Huang, X.; Zhang, S.; Elgarnal, A.; and Metaxas, D. 2016a. Spda-cnn: Unifying semantic part detection and abstraction for fine-grained recognition. *Proc. CVPR*.
- Zhang, J.; Lin, Z.; Brandt, J.; Shen, X.; and Sclaroff, S. 2016b. Top-down neural attention by excitation backprop. *arXiv:1608.00507*.
- Zhao, B.; Wu, X.; Feng, J.; Peng, Q.; and Yan, S. 2016. Diversified visual attention networks for fine-grained object classification. *arXiv:1606.08572*.