



Unsupervised object discovery and co-localization by deep descriptor transformation

Xiu-Shen Wei^{a,1}, Chen-Lin Zhang^{b,1}, Jianxin Wu^{b,*}, Chunhua Shen^{c,*}, Zhi-Hua Zhou^b

^a Megvii Research Nanjing, Megvii Technology Ltd. (Face++), China

^b National Key Laboratory for Novel Software Technology, Nanjing University, China

^c The University of Adelaide, Adelaide, Australia

ARTICLE INFO

Article history:

Received 26 March 2018

Revised 9 October 2018

Accepted 21 October 2018

Available online 7 November 2018

Keywords:

Unsupervised object discovery

Object co-localization

Deep descriptor transformation

Pre-trained CNN models

ABSTRACT

Reusable model design becomes desirable with the rapid expansion of computer vision and pattern recognition applications. In this paper, we focus on the reusability of pre-trained deep convolutional models. Specifically, different from treating pre-trained models as feature extractors, we reveal more treasures beneath convolutional layers, i.e., the convolutional activations could act as a detector for the common object in the object co-localization problem. We propose a simple yet effective method, termed Deep Descriptor Transformation (DDT), for evaluating the correlations of descriptors and then obtaining the category-consistent regions, which can accurately locate the common object in a set of unlabeled images, i.e., object co-localization. Empirical studies validate the effectiveness of the proposed DDT method. On benchmark object co-localization datasets, DDT consistently outperforms existing state-of-the-art methods by a large margin. Moreover, DDT also demonstrates good generalization ability for unseen categories and robustness for dealing with noisy data. Beyond those, DDT can be also employed for harvesting web images into valid external data sources for improving performance of both image recognition and object detection.

© 2018 Published by Elsevier Ltd.

1. Introduction

Model reuse [1] attempts to construct a model by utilizing existing available models, mostly trained for other tasks, rather than building a model from scratch. Particularly in deep learning, since deep convolutional neural networks have achieved great success in various tasks involving images, videos, texts and more, there are several studies have the flavor of reusing deep models pre-trained on *ImageNet* [2].

In pattern recognition and computer vision, pre-trained models on *ImageNet* have been successfully adopted to various usages, e.g., as universal feature extractors [3–6], object proposal generators [7], etc. In particular, [8] proposed the SCDA (Selective Convolutional Descriptor Aggregation) method to utilize pre-trained models for both localizing a single fine-grained object (e.g., birds of different species) in each image and retrieving fine-grained images of the same classes/species in an unsupervised fashion.

In this paper, we reveal that the convolutional activations can be used as a detector for the *common object* in object co-localization. Object co-localization is a fundamental computer vision problem, which simultaneously localizes objects of the same category across a set of distinct images. Specifically, we propose a simple but effective method termed Deep Descriptor Transformation (DDT) for object co-localization. In DDT, the deep convolutional descriptors extracted from pre-trained deep convolutional models are transformed into a new space, where it can evaluate the correlations between these descriptors. By leveraging the correlations among images in the image set, the common object inside these images can be located automatically without additional supervision signals. The pipeline of DDT is shown in Fig. 1. In general, our DDT could demonstrate the possibility of convolutional activations/descriptors in pre-trained models *being able to act as a detector for the common object*.

Experimental results show that DDT significantly outperforms existing state-of-the-art methods, including object co-localization and weakly supervised object localization, in both the deep learning and hand-crafted feature scenarios. Besides, we empirically show that DDT has a good generalization ability for unseen images apart from *ImageNet*. More importantly, the proposed method is robust, because DDT can also detect the noisy images which

* Corresponding authors.

E-mail addresses: weixiushen@megvii.com (X.-S. Wei), wujx2001@gmail.com (J. Wu), chunhua.shen@adelaide.edu.au (C. Shen).

¹ The first two authors contributed equally to this work.

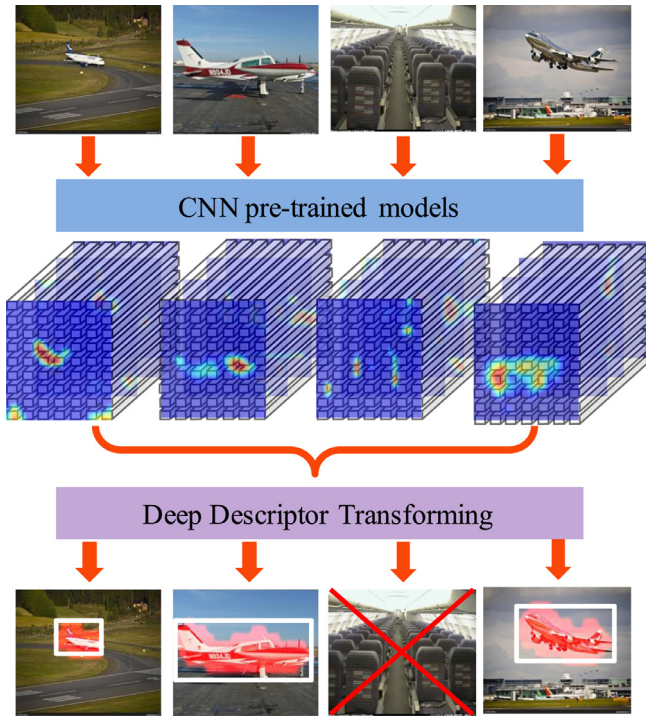


Fig. 1. Pipeline of the proposed DDT method for object co-localization. In this instance, the goal is to localize the *airplane* within each image. Note that, there might be few noisy images in the image set. (Best viewed in color.).

do not contain the common object. Thanks to the advantages of DDT, our method could be used as a tool to harvest easy-to-obtain but noisy web images. We can employ DDT to remove noisy images from webly datasets for improving image recognition accuracy. Moreover, it can be also utilized to supply object bounding boxes of web images. Then, we use these images with automatically labeled object boxes as valid external data sources to enhance object detection performance.

Our main contributions are as follows:

1. We propose a simple yet effective method, i.e., Deep Descriptor Transformation, for unsupervised object discovery and object co-localization. Besides, DDT reveals another probability of deep pre-trained network reusing, i.e., convolutional activations/descriptors can play a role as a common object detector.
2. The co-localization process of DDT is both effective and efficient, which does not require image labels, negative images or redundant object proposals. DDT consistently outperforms state-of-the-arts of object co-localization methods and weakly supervised object localization methods. With the ensemble of multiple CNN layers, DDT could further improve its co-localization performance.
3. DDT has a good generalization ability for unseen categories and robustness for dealing with noisy data. Thanks to these advantages, DDT can be employed beyond the narrow co-localization task. Specifically, it can be used as a generalized tool for exploiting noisy but free web images. By removing noisy images and automatically supplying object bounding boxes, these web images processed by DDT could become valid external data sources for improving both recognition and detection performance. We thus provide a very useful tool for automatically annotating images. The effectiveness of DDT augmentation on recognition and detection is validated in Section 4.6.

4. Based on the previous point, we also collect an object detection dataset from web images, named *WebVOC*.² It shares the same 20 categories as the *PASCAL VOC* dataset [9], and has a similar dataset scale (10k images) comparing with *PASCAL VOC*. We also release the *WebVOC* dataset with the automatically generated bounding boxes by DDT for further study.

This paper is extended based on our preliminary work [10]. Comparing with it, we now further introduce the multiple layer ensemble strategy for improving co-localization performance, provide DDT augmentation for handling web images, apply the proposed method on webly-supervised learning tasks (i.e., both recognition and detection), and supply our DDT based webly object detection dataset.

The remainder of the paper is organized as follows. In Section 2, we briefly review related literature of CNN model reuse, object co-localization and webly-supervised learning. In Section 3, we introduce our proposed method (DDT and its variant DDT⁺). Section 4 reports the object co-localization results and the results of webly-supervised learning tasks. We conclude the paper in Section 5 finally.

2. Related work

We briefly review three lines of related work: model reuse of CNNs, research on object co-localization and webly-supervised learning.

2.1. CNN model reuse

Reusability has been emphasized by [1] as a crucial characteristic of the new concept of *learnware*. It would be ideal if models can be reused in scenarios that are very different from their original training scenarios. Particularly, with the breakthrough in image classification using Convolutional Neural Networks (CNN), pre-trained CNN models trained for one task (e.g., recognition) have also been applied to domains different from their original purposes (e.g., for describing texture [5] or finding object proposals [7]). Meanwhile, several works reveal that deep CNN models designed for classification can emerge object detectors [11,12]. However, for such adaptations of pre-trained models, they still require further annotations in the new domain (e.g., image labels). While, DDT deals with the object co-localization problem in a setting whose only assumption is these images contain one common object, rather than requiring image labels.

Coincidentally, several recent works also shed lights on CNN pre-trained model reuse in the unsupervised setting, e.g., SCDA (Selective Convolutional Descriptor Aggregation) [8]. SCDA is proposed for handling the fine-grained image retrieval task, where it uses pre-trained models (from *ImageNet*) to locate main objects in fine-grained images. It is the most related work to ours, even though SCDA is not for object co-localization. Different from our DDT, SCDA assumes only an object of interest in each image, and meanwhile objects from other categories does not exist. Thus, SCDA locates the object using cues from this *single image* assumption. Clearly, it can not work well for images containing diverse objects (cf. Table 2 and Table 3), and also can not handle data noise (cf. Section 4.5).

2.2. Object co-localization

Object co-localization [13–16], is a fundamental problem in pattern recognition and computer vision, where it needs to dis-

² Our *WebVOC* dataset and the source codes of DDT are both released via: <http://lamda.nju.edu.cn/weixs/project/DDT/DDT.html>.

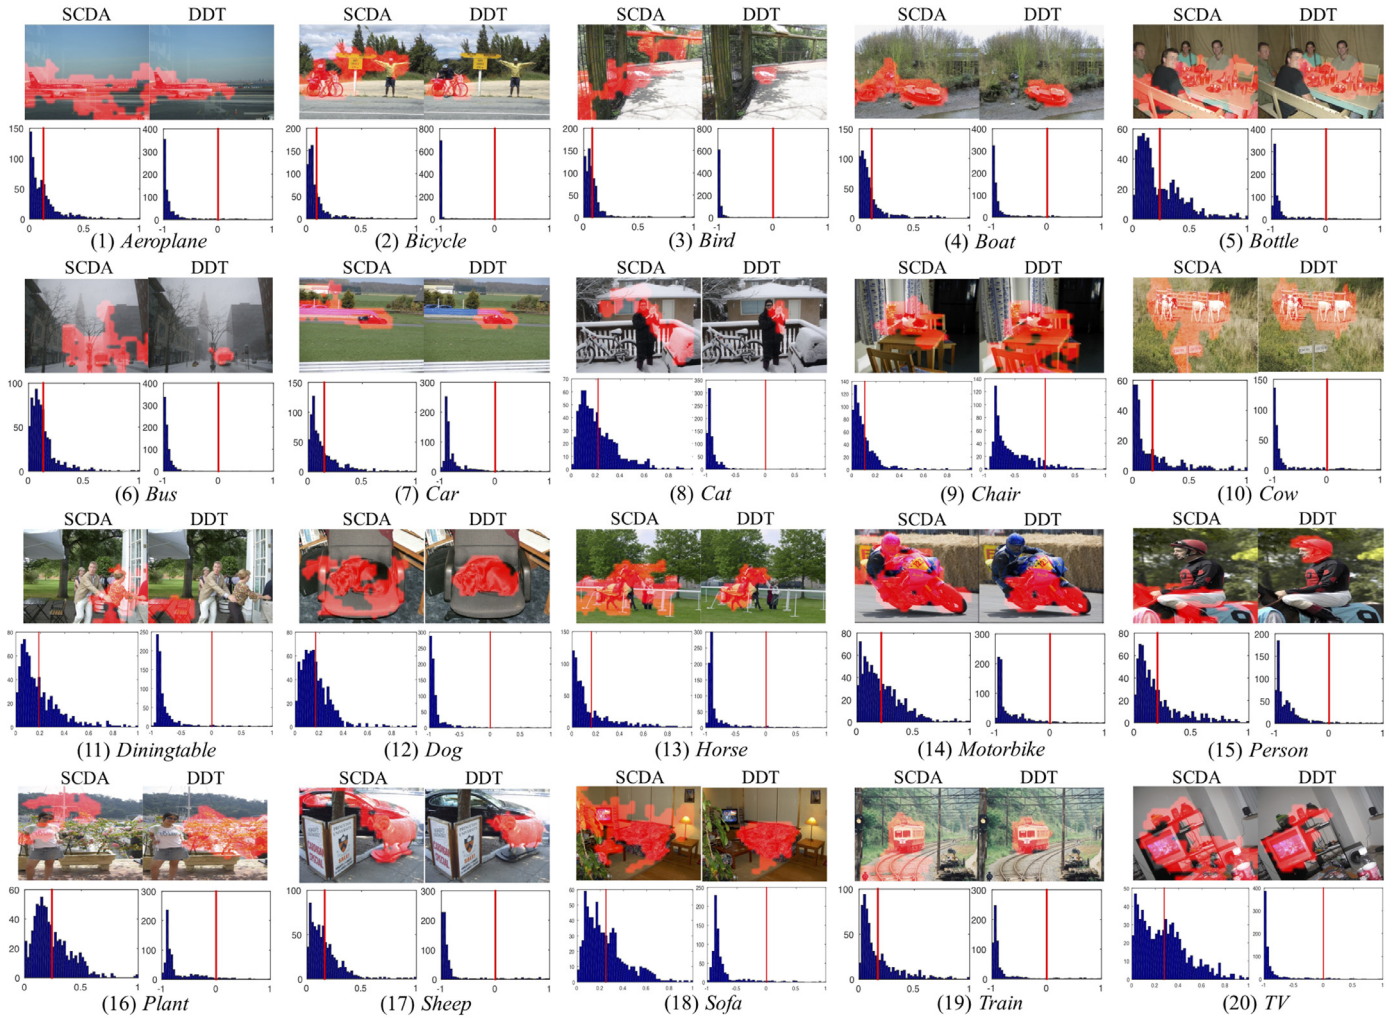


Fig. 2. Examples of twenty categories from the PASCAL VOC 2007 dataset [9]. The first column of each sub-figure is produced by SCDA, the second column is by our DDT. The red vertical lines in the histogram plots indicate the corresponding thresholds for localizing objects. The selected regions in images are highlighted in red. (Best viewed in color and zoomed in.). (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

cover the common object emerging in only positive sets of example images (without any negative examples or further supervisions). Object co-localization shares some similarities with image co-segmentation [17–19]. Instead of generating a precise segmentation of the related objects in each image, co-localization methods aim to return a bounding box around the object. Moreover, it also allows us to extract rich features from within the boxes to compare across images, which has shown to be very helpful for detection [20].

Additionally, co-localization is also related to weakly supervised object localization (WSOL) [21–25]. But the key difference between them is that WSOL requires manually-labeled negative images whereas co-localization does not. Thus, WSOL methods could achieve better localization performance than co-localization methods. However, our proposed methods perform comparably with most state-of-the-art WSOL methods (cf. Table 4).

In the literature, some representative co-localization methods are based on low-level visual cues and optimization algorithms [20]. formulates co-localization as a boolean constrained quadratic program which can be relaxed to a convex problem. Then, it is further accelerated by the Frank-Wolfe algorithm [26]. After that, [14] proposes a Probabilistic Hough Matching algorithm to match object proposals across images and then dominant objects are localized by selecting proposals based on matching scores.

Recently, there also emerge several co-localization methods based on pre-trained deep convolutional models, e.g., [6,24]. Unfortunately, these methods just treat pre-trained models as simple feature extractors to extract the fully connected representations, which do not sufficiently mine the treasures beneath the convolutional layers (i.e., leveraging the original correlations between deep descriptors among convolutional layers). Moreover, these methods also require object proposals as a part of their object discovery, which not only made them highly depend on the quality of object proposals, but may lead to huge computational costs. In addition, many previous co-localization methods can not handle noisy data, except for [20].

Compared to the previous work, DDT works without requiring bounding boxes, additional image labels or redundant object proposals. Images only need one forward run through a pre-trained model. Then, efficient deep descriptor transformation is employed for obtaining the category-consistent image regions. DDT is very easy to implement, and surprisingly has good generalization ability and robustness. Furthermore, DDT can be used a valid data augmentation tool for handling noisy but free web images.

2.3. Webly-supervised learning

Recent development of deep CNNs has led to great success in a variety of computer vision tasks. This success is largely driven by




Fig. 3. Examples from five randomly sampled categories of PASCAL VOC 2007 [9]. The red highlighted regions in images are detected as containing common objects by our proposed methods. In each sub-figure, the first column presents the prediction by our DDT (cf. Algorithm 2). The middle column shows the DDT's result based on the lower convolutional layer. The last column are the predicted results by our DDT+. (Best viewed in color and zoomed in.) (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

the availability of large scale well-annotated image datasets, e.g., ImageNet [2], MS COCO [27] and PASCAL VOC [9]. However, annotating a massive number of images is extremely labor-intensive and costly. To reduce the annotation labor costs, an alternative approach is to obtain the image annotations directly from the image search engine from the Internet, e.g., Google or Bing.

However, the annotations of web images returned by a search engine will inevitably be noisy since the query keywords may not be consistent with the visual content of target images. Thus, webly-supervised learning methods [28,29] are proposed for overcoming this issue.

There are two main branches of webly-supervised learning. The first branch attempts to boost existing object recognition task performance using web resources [30–32]. Some works are implemented as semi-supervised frameworks by first generating a small group of labeled seed images and then enlarging the dataset from these seeds via web data, e.g., [31,32]. In very recently, [30] proposes a two-level attention framework for dealing with webly-supervised classification, which achieves a new state-of-the-art. Specifically, they not only use a high-level attention focusing on

a group of images for filtering out noisy images, but also employ a low-level attention for capturing the discriminative image regions on the single image level

The second branch is learning visual concepts directly from the web, e.g., [33,34]. Methods belonging to this category usually collected a large image pool from image search engines and then performed a filtering operation to remove noise and discover visual concepts. Our strategy for handling web data based on DDT naturally falls into the second category. In practice, since DDT could (1) recognize noisy images and also (2) supply bounding boxes of objects, we leverage the first usage of DDT to handle webly-supervised classification (cf. Table 6 and Table 7), and leverage both two usages to deal with webly-supervised detection (cf. Table 8 and Table 9).

3. The proposed method

In this section, we propose the Deep Descriptor Transformation (DDT) method. Firstly, we introduce notations used in this paper. Then, we present the DDT process followed by discussions and analyses. Finally, in order to further improve the object co-localization performance, the multiple layer ensemble strategy is utilized in DDT.

3.1. Preliminaries

The following notations are used in the rest of this paper. The term “feature map” indicates the convolution results of one channel; the term “activations” indicates feature maps of all channels in a convolution layer; and the term “descriptor” indicates the d -dimensional component vector of activations.

Given an input image I of size $H \times W$, the activations of a convolution layer are formulated as an order-3 tensor T with $h \times w \times d$ elements. T can be considered as having $h \times w$ cells and each cell contains one d -dimensional deep descriptor. For the n -th image in the image set, we denote its corresponding deep descriptors as $X^n = \{\mathbf{x}_{(i,j)}^n \in \mathcal{R}^d\}$, where (i, j) is a particular cell ($i \in \{1, \dots, h\}$, $j \in \{1, \dots, w\}$) and $n \in \{1, \dots, N\}$.

3.2. SCDA recap

Since SCDA (Selective Convolutional Descriptor Aggregation) [8] is the most related work to ours, we hereby present a recap of this method. SCDA is proposed for dealing with the fine-grained image retrieval problem. It employs pre-trained CNN models to select the meaningful deep descriptors by localizing the main object in fine-grained images unsupervisedly. In SCDA, it assumes that each image contains only one main object of interest and without other categories' objects. Thus, the object localization strategy is based on the activation tensor of a *single* image.

Concretely, for an image, the activation tensor is added up through the depth direction. Thus, the $h \times w \times d$ 3-D tensor becomes a $h \times w$ 2-D matrix, which is called the “aggregation map” in SCDA. Then, the mean value \bar{a} of the aggregation map is regarded as the threshold for localizing the object. If the activation response in the position (i, j) of the aggregation map is larger than \bar{a} , it indicates the object might appear in that position.

3.3. Deep descriptor transformation (DDT)

What distinguishes DDT from SCDA is that we can leverage the correlations beneath the whole *image set*, instead of a *single* image. Additionally, different from weakly supervised object localization, we do not have either image labels or negative image sets in

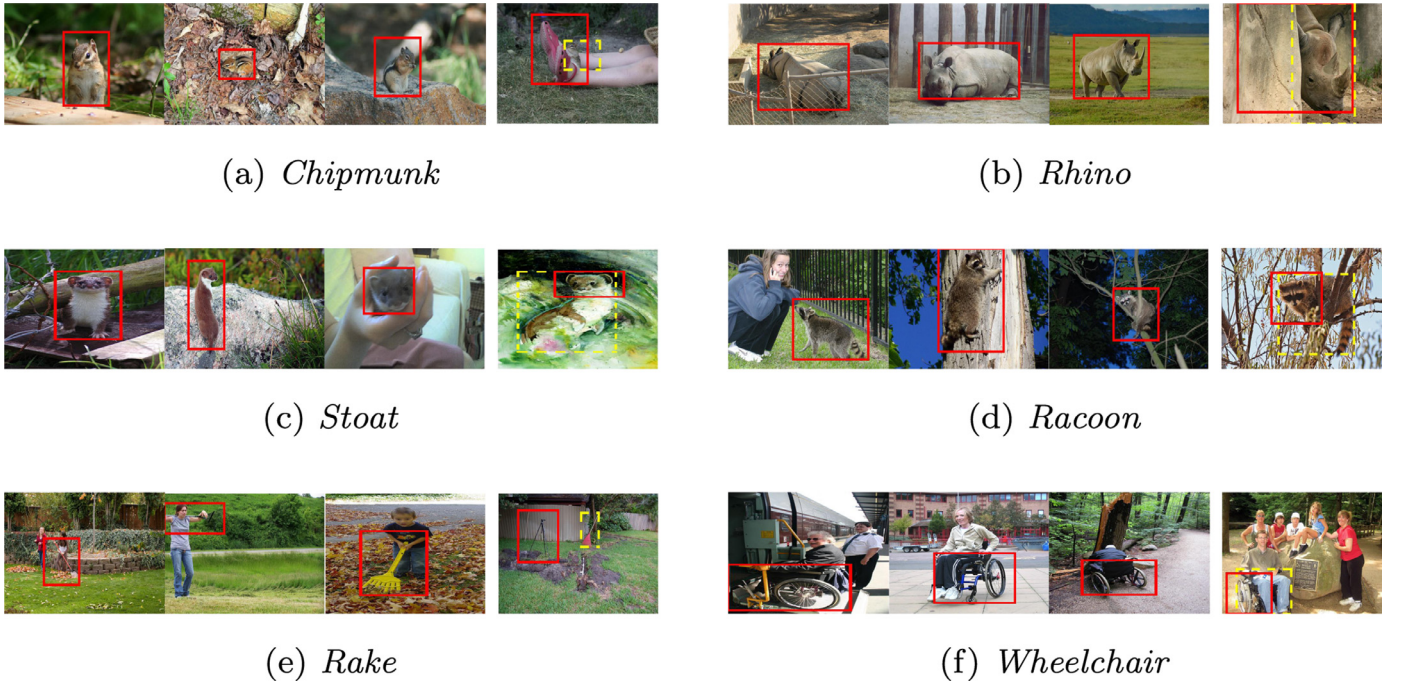


Fig. 4. Random samples of predicted object co-localization bounding box on *ImageNet Subsets*. Each sub-figure contains three successful predictions and one failure case. In these images, the red rectangle is the prediction by DDT, and the yellow dashed rectangle is the ground truth bounding box. In the successful predictions, the yellow rectangles are omitted since they are exactly the same as the red predictions. (Best viewed in color and zoomed in.). (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)




Fig. 5. ROC curves illustrating the effectiveness of our DDT at identifying noisy images on the *Object Discovery* dataset. The curves in red line are the ROC curves of DDT. The curves in blue dashed line present the method in [20]. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

WSOL, so that the information we can use is only from the pre-trained models. Here, we transform the deep descriptors in convolutional layers to mine the hidden cues for co-localizing common objects.

Principal component analysis (PCA) [35] is a statistical procedure, which uses an orthogonal transformation to convert a set of observations of possibly correlated variables into a set of linearly uncorrelated variables (i.e., the principal components). This transformation is defined in such a way that the first principal component has the largest possible variance, and each succeeding component in turn has the highest variance possible under the constraint that it is orthogonal to all the preceding components.

PCA is widely used in computer vision and pattern recognition for image denoising, 3D object retrieval, statistical shape modeling, subspace learning, and so on. Specifically, in this paper, we uti-

lize PCA as projection directions for transforming these deep descriptors $\{\mathbf{x}_{(i,j)}\}$ to evaluate their correlations. Then, on each projection direction, the corresponding principal component's values are treated as the cues for object co-localization, especially the first principal component. Thanks to the property of this kind of transforming, DDT is also able to handle data noise.

In DDT, for a set of N images containing objects from the same category, we first collect the corresponding convolutional descriptors (X^1, \dots, X^N) from the last convolutional layer by feeding the images into a pre-trained CNN model. Then, the mean vector of all the descriptors is calculated by:

$$\bar{\mathbf{x}} = \frac{1}{K} \sum_n \sum_{i,j} \mathbf{x}_{(i,j)}^n, \quad (1)$$

where $K = h \times w \times N$. Note that, here we assume each image has the same number of deep descriptors (i.e., $h \times w$) for presentation clarity. Our proposed method, however, can handle input images with arbitrary resolutions.

Then, after obtaining the covariance matrix:

$$\text{Cov}(\mathbf{x}) = \frac{1}{K} \sum_n \sum_{i,j} (\mathbf{x}_{(i,j)}^n - \bar{\mathbf{x}})(\mathbf{x}_{(i,j)}^n - \bar{\mathbf{x}})^\top, \quad (2)$$

we can get the eigenvectors ξ_1, \dots, ξ_d of $\text{Cov}(\mathbf{x})$ which correspond to the sorted eigenvalues $\lambda_1 \geq \dots \geq \lambda_d \geq 0$.

As aforementioned, since the first principal component has the largest variance, we take the eigenvector ξ_1 corresponding to the largest eigenvalue as the main projection direction. For the deep descriptor at a particular position (i, j) of an image, its first principal component p^1 is calculated as follows:

$$p_{(i,j)}^1 = \xi_1^\top (\mathbf{x}_{(i,j)} - \bar{\mathbf{x}}). \quad (3)$$

According to their spatial locations, all $p_{(i,j)}^1$ from an image are formed into a 2-D matrix whose dimensions are $h \times w$. We call that matrix as **indicator matrix**:

$$P^1 = \begin{bmatrix} p_{(1,1)}^1 & p_{(1,2)}^1 & \dots & p_{(1,w)}^1 \\ p_{(2,1)}^1 & p_{(2,2)}^1 & \dots & p_{(2,w)}^1 \\ \vdots & \vdots & \ddots & \vdots \\ p_{(h,1)}^1 & p_{(h,2)}^1 & \dots & p_{(h,w)}^1 \end{bmatrix} \in \mathcal{R}^{h \times w}. \quad (4)$$

P^1 contains positive (negative) values which can reflect the positive (negative) correlations of these deep descriptors. The larger the absolute value is, the higher the positive (negative) correlation will be. Because ξ_1 is obtained through all N images in that image set, the positive correlation could indicate the *common characteristic* through N images. Specifically, in the object co-localization scenario, the corresponding positive correlation indicates indeed the *common object* inside these images.

Therefore, the value zero could be used as a natural threshold for dividing P^1 of one image into two parts: one part has positive values indicating the common object, and the other part has negative values presenting background or objects that rarely appear. Additionally, if P^1 of an image has no positive value, it indicates that no common object exists in that image, which can be used for detecting noisy images.

In practice, for localizing objects, P^1 is resized by the nearest interpolation, such that its size is the same as that of the input image. Since the nearest interpolation is the zero-order interpolation method, it will not change the signs of the numbers in P^1 . Thus, the resized P^1 can be used for localizing the common object according to the aforementioned principle with the natural threshold (i.e., the value zero). Meanwhile, we employ the algorithm described in Algorithm 1 to collect the largest connected component of the positive regions in the resized P^1 to remove several small noisy positive parts. Then, the minimum rectangle bounding box which contains the largest connected component of positive regions is returned as our object co-localization prediction for each image. The whole procedure of the proposed DDT method is shown in Algorithm 2.

3.4. Discussions and analyses

In this section, we investigate the effectiveness of DDT by comparing with SCDA.

As shown in Fig. 2, the object localization regions of SCDA and DDT are highlighted in red. Because SCDA only considers the information from a single image, for example, in Fig. 2 (2), “bike”, “person” and even “guide-board” are all detected as main objects. Similar observations could be found in Fig. 2 (5), (13), (17), (18), etc.

Algorithm 1 Finding the largest connected component.

Require: The resized indicator matrix P^1 corresponding to an image I ;

- 1: Transform P^1 into a binary map \hat{P}^1 ,
where $\hat{p}_{(i,j)}^1 = \begin{cases} 1 & \text{if } p_{(i,j)}^1 > 0 \\ 0 & \text{otherwise} \end{cases}$;
- 2: Select one pixel p in \hat{P}^1 as the starting point;
- 3: **while** True **do**
- 4: Use a flood-fill algorithm to label all the pixels in the connected component containing p ;
- 5: **if** All the pixels are labeled **then**
- 6: Break;
- 7: **end if**
- 8: Search for the next unlabeled pixel as p ;
- 9: **end while**
- 10: Obtain the connectivity of the connected components, and their corresponding size (pixel numbers);
- 11: Select the connected component \hat{P}_c^1 with the largest pixel number;
- 12: **return** The largest connected component \hat{P}_c^1 .

Algorithm 2 Deep descriptor transformation (DDT).

Require: A set of N images containing the common object, and a pre-trained CNN model \mathcal{F} ;

- 1: Feed these images with their original resolutions into \mathcal{F} ;
- 2: Collect the corresponding convolutional descriptors X^1, \dots, X^N from the last convolutional layer of \mathcal{F} ;
- 3: Calculate the mean vector $\bar{\mathbf{x}}$ of all the descriptors using Eq. 1;
- 4: Compute the covariance matrix $\text{Cov}(\mathbf{x})$ of these deep descriptors based on Eq. 2;
- 5: Compute the eigenvectors ξ_1, \dots, ξ_d of $\text{Cov}(\mathbf{x})$;
- 6: Select ξ_1 with the largest eigenvalue as the main transformation direction;
- 7: **repeat**
- 8: Calculate the indicator matrix P^1 for image I based on Eq. 2 and Eq. 3;
- 9: Resize P^1 into its image's resolution by nearest interpolation;
- 10: Collect the largest connected component \hat{P}_c^1 of these positive regions of the resized P^1 by Algo. 4;
- 11: Obtain the minimum rectangle bounding box covering \hat{P}_c^1 as the prediction;
- 12: **until** All the N images are done;
- 13: **return** The minimum rectangle bounding boxes.

Furthermore, we normalize the values (all positive) of the aggregation map of SCDA into the scale of $[0, 1]$, and calculate the mean value (which is taken as the object localization threshold in SCDA). The histogram of the normalized values in aggregation map is also shown in the corresponding sub-figure in Fig. 2. The red vertical line corresponds to the threshold. We can find that, beyond the threshold, there are still many values. It gives an explanation about why SCDA highlights more regions.

Whilst, for DDT, it leverages the whole image set to transform these deep descriptors into P^1 . Thus, for the *bicycle* class (cf. Fig. 2 (2)), DDT can accurately locate the “bicycle” object. The histogram of DDT is also drawn. But, P^1 has both positive and negative values. We normalize P^1 into the $[-1, 1]$ scale this time. Apparently, few values are larger than the DDT threshold (i.e., the value zero). More importantly, many values are close to -1 which indicates the strong negative correlation. This observation validates the effectiveness of DDT in object co-localization. As another example shown in Fig. 2 (11), SCDA even wrongly locates “person” in the



Fig. 6. Examples of noisy images in the WebCars dataset recognized by DDT.

image belonging to the *diningtable* class. While, DDT can correctly and accurately locate the “diningtable” image region. More examples are presented in Fig. 2. In that figure, some failure cases can be also found, e.g., the *chair* class in Fig. 2 (9).

In addition, the normalized P^1 can be also used as localization probability scores. Combining it with conditional random field techniques might produce more accurate object boundaries. Thus, DDT can be modified slightly in that way, and then perform the co-segmentation problem.

3.5. Multiple layer ensemble

As is well known, CNNs are composed of multiple processing layers to learn representations of images with multiple levels of abstraction. Different layers will learn different level visual information [36]. Lower layers have more general representations (e.g., textures and shapes), and they can capture more detailed visual cues. By contrast, the learned representations of deeper layers contain more semantic information (i.e., high-level concepts). Thus, deeper layers are good at abstraction, but they lack visual details. Apparently, lower layer and deeper layer are complementary with each other. Based on this, several previous work, e.g., [37,38], aggregate the information of multiple layers to boost the final performance on their computer vision tasks.

Inspired by them, we also incorporate the lower convolutional layer in pre-trained CNNs to supply finer detailed information for object co-localization, which is named as DDT⁺.

Concretely, as aforementioned in Algorithm 2, we can obtain \hat{P}_c^1 of the resize P^1 for each image from the last convolutional layer by our DDT. Several visualization examples of \hat{P}_c^1 are shown in the first column of Fig. 3. In DDT⁺, beyond that, those deep descriptors from the previous convolutional layer before the last one are also used for generating its corresponding resized P^1 , which is notated as P_{prev}^1 . For P_{prev}^1 , we directly transform it into a binary map \hat{P}_{prev}^1 . In the middle column of Fig. 3, the red highlighted regions represent the co-localization results by \hat{P}_{prev}^1 . Since the activations from the previous convolutional layer are less related to the high-level semantic meaning than those from the last convolutional layer, other objects not belonging to the common object category are also being detected. However, the localization boundaries are much finer than \hat{P}_c^1 . Therefore, we combine \hat{P}_c^1 and \hat{P}_{prev}^1 together to obtain the final co-localization prediction as follows:

$$\hat{P}_c^1 \cap \hat{P}_{prev}^1. \quad (5)$$

As shown in the last column of Fig. 3, the co-localization visualization results of DDT⁺ are better than the results of DDT, especially for the *bottle* class. In addition, from the quantitative perspective, DDT⁺ will bring on average 1.5% improvements on object co-localization (cf. Table 2, Table 3 and Table 5).

4. Experiments

In this section, we first introduce the evaluation metric and datasets used in object co-localization. Then, we compare the empirical results of our DDT and DDT⁺ with other state-of-the-arts on these datasets. The computational cost is reported too. Moreover, the results in Section 4.4 and Section 4.5 illustrate the generalization ability and robustness of the proposed method. Furthermore, we will discuss the ability of DDT to utilize web data as valid augmentation for improving the accuracy of traditional image recognition and object detection tasks. Finally, the further study in Section 4.7 reveals DDT might deal with part-based object co-localization, which is a novel and challenging problem.

In our experiments, the images keep the original image resolutions. For the pre-trained deep model, the publicly available VGG-19 model [39] is employed to perform DDT by extracting deep convolution descriptors from the last convolution layer (i.e., the $relu_{5,4}$ layer) and employed to perform DDT⁺ by using both the last convolution layer (i.e., the $relu_{5,4}$ layer) and its previous layer (i.e., the $relu_{5,3}$ layer). We use the open-source library MatConvNet [40] for conducting experiments. All the experiments are run on a computer with Intel Xeon E5-2660 v3, 500G main memory, and a K80 GPU.

4.1. Evaluation metric and datasets

Following previous object co-localization works [6,14,20], we take the correct localization (CorLoc) metric for evaluating the proposed method. CorLoc is defined as the percentage of images correctly localized according to the PASCAL-criterion [9]:

$$\frac{\text{area}(B_p \cap B_{gt})}{\text{area}(B_p \cup B_{gt})} > 0.5, \quad (6)$$

where B_p is the predicted bounding box and B_{gt} is the ground-truth bounding box. All CorLoc results are reported in percentages.

Our experiments are conducted on four challenging datasets commonly used in object co-localization, i.e., the *Object Discovery* dataset [41], the *PASCAL VOC 2007/VOC 2012* dataset [9] and the *ImageNet Subsets* dataset [6].

For experiments on the *PASCAL VOC* datasets, we follow [6,14,26] to use all images in the *trainval* set (excluding images that only contain object instances annotated as *difficult* or *truncated*). For *Object Discovery*, we use the 100-image subset following [14,41] in order to make an appropriate comparison with other methods.

In addition, *Object Discovery* has 18%, 11% and 7% noisy images in the *Airplane*, *Car* and *Horse* categories, respectively. These noisy images contain no object belonging to their category, as the third image shown in Fig. 1. Particularly, in Section 4.5, we quantitatively measure the ability of our proposed DDT to identify these noisy images.

To further investigate the generalization ability of DDT, *ImageNet Subsets* [6] are used, which contain six subsets/categories. These subsets are held-out categories from the 1000-label ILSVRC classification [2]. That is to say, these subsets are “unseen” by pre-trained CNN models. Experimental results in Section 4.4 show that our proposed methods is insensitive to the object category.

4.2. Comparisons with state-of-the-arts

In this section, we compare the object co-localization performance of our methods with state-of-the-art methods including both object co-localization and weakly supervised object localization.



Fig. 7. Examples of our WebVOC detection dataset. The red bounding boxes in these figures are automatically labeled by the proposed DDT method. (Best viewed in color and zoomed in.). (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

Table 1
Comparisons of CorLoc on *object discovery*.

Methods	airplane	car	horse	Mean
Joulin et al. [42]	32.93	66.29	54.84	51.35
Joulin et al. [19]	57.32	64.04	52.69	58.02
Rubinstein et al. [41]	74.39	87.64	63.44	75.16
Tang et al. [20]	71.95	93.26	64.52	76.58
SCDA [8]	87.80	86.52	75.37	83.20
Cho et al. [14]	82.93	94.38	75.27	84.19
Our DDT	91.46	95.51	77.42	88.13
Our DDT ⁺	91.46	94.38	76.34	87.39

4.2.1. Comparisons to object co-localization methods

We first compare the results of DDT to state-of-the-arts (including SCDA) on *Object Discovery* in Table 1. For SCDA, we also use VGG-19 to extract the convolution descriptors and perform experiments. As shown in that table, DDT outperforms other methods by about 4% in the mean CorLoc metric. Especially for the *airplane* class, it is about 10% higher than that of [14]. In addition, note that the images of each category in this dataset contain only one object, thus, SCDA can perform well. But, our DDT⁺ gets a slightly lower CorLoc score than DDT, which is an exception in all the object co-localization datasets. In fact, for *car* and *horse* of the *Object Discovery* dataset, DDT⁺ only returns one more wrong prediction than DDT for each category.

For *PASCAL VOC 2007* and *2012*, these datasets contain diverse objects per image, which is more challenging than *Object Discovery*. The comparisons of the CorLoc metric on these two datasets are reported in Table 2 and Table 3, respectively. It is clear that on average our DDT and DDT⁺ outperform the previous state-of-the-arts (based on deep learning) by a large margin on both datasets. Moreover, our methods work well on localizing small common objects, e.g., “bottle” and “chair”. In addition, because most images of these datasets have multiple objects, which do not obey SCDA’s assumption, SCDA performs poorly in the complicated environment. For fair comparisons, we also use VGG-19 to extract the fully connected representations of the object proposals in [6], and then perform the remaining processes of their method (the source codes are provided by the authors). As aforementioned, due to the high

dependence on the quality of object proposals, their mean CorLoc metric of VGG-19 is 41.9% and 45.6% on *VOC 2007* and *2012*, respectively. The improvements are limited, and the performance is still significantly worse than ours.

4.2.2. Comparisons to weakly supervised localization methods

To further verify the effectiveness of our methods, we also compare DDT and DDT⁺ with some state-of-the-art methods for weakly supervised object localization. Table 4 illustrates these empirical results on *VOC 2007*. Particularly, DDT achieves 46.9% on average which is higher than most WSOL methods in the literature, and DDT⁺ achieves 48.5% on average. While, our proposed method still has a gap comparing with state-of-the-art WSOL methods, e.g., [43,44]. But, it is understandable as our methods do *not* use any negative data for co-localization. In addition, our methods could handle noisy data (cf. Section 4.5). However, existing WSOL methods are not designed to deal with noise.

4.3. Computational costs of DDT/DDT⁺

Here, we take the total 171 images in the *aeroplane* category of *VOC 2007* as examples to report the computational costs. The average image resolution of the 171 images is 350×498 . The computational time of DDT has two main components: one is for feature extraction, the other is for deep descriptor transformation (cf. Algorithm 2). Because we just need the first principal component, the transformation time on all the 120,941 descriptors of 512-d is only 5.7 seconds. The average descriptor extraction time is 0.18 second/image on GPU and 0.86 second/image on CPU, respectively. For DDT⁺, it has the same deep descriptor extraction time. Although it needs descriptors from two convolutional layers, it only requires one time feed-forward processing. The deep descriptor transformation time of DDT⁺ is only 11.9 seconds for these 171 images. These numbers above could ensure the efficiency of the proposed methods in real-world applications.

4.4. Unseen classes apart from ImageNet

In order to justify the generalization ability of the proposed methods, we also conduct experiments on some images (of six

Table 2

Comparisons of the CorLoc metric with state-of-the-art co-localization methods on VOC 2007.

Methods	aero	bike	bird	boat	bottle	bus	car	cat	chair	cow	table*	dog	horse	mbike	person	plant	sheep	sofa	train	tv	Mean
Joulin et al. [26]	32.8	17.3	20.9	18.2	4.5	26.9	32.7	41.0	5.8	29.1	34.5	31.6	26.1	40.4	17.9	11.8	25.0	27.5	35.6	12.1	24.6
SCDA [8]	54.4	27.2	43.4	13.5	2.8	39.3	44.5	48.0	6.2	32.0	16.3	49.8	51.5	49.7	7.7	6.1	22.1	22.6	46.4	6.1	29.5
Cho et al. [14]	50.3	42.8	30.0	18.5	4.0	62.3	64.5	42.5	8.6	49.0	12.2	44.0	64.1	57.2	15.3	9.4	30.9	34.0	61.6	31.5	36.6
Li et al. [6]	73.1	45.0	43.4	27.7	6.8	53.3	58.3	45.0	6.2	48.0	14.3	47.3	69.4	66.8	24.3	12.8	51.5	25.5	65.2	16.8	40.0
Our DDT	67.3	63.3	61.3	22.7	8.5	64.8	57.0	80.5	9.4	49.0	22.5	72.6	73.8	69.0	7.2	15.0	35.3	54.7	75.0	29.4	46.9
Our DDT ⁺	71.4	65.6	64.6	25.5	8.5	64.8	61.3	80.5	10.3	49.0	26.5	72.6	75.2	69.0	9.9	12.2	39.7	55.7	75.0	32.5	48.5

Table 3

Comparisons of the CorLoc metric with state-of-the-art co-localization methods on VOC 2012.

Methods	aero	bike	bird	boat	bottle	bus	car	cat	chair	cow	table*	dog	horse	mbike	person	plant	sheep	sofa	train	tv	Mean
SCDA [8]	60.8	41.7	38.6	21.8	7.4	67.6	38.8	57.4	16.0	34.0	23.9	53.8	47.3	54.8	7.9	9.9	25.3	23.2	50.2	10.1	34.5
Cho et al. [14]	57.0	41.2	36.0	26.9	5.0	81.1	54.6	50.9	18.2	54.0	31.2	44.9	61.8	48.0	13.0	11.7	51.4	45.3	64.6	39.2	41.8
Li et al. [6]	65.7	57.8	47.9	28.9	6.0	74.9	48.4	48.4	14.6	54.4	23.9	50.2	69.9	68.4	24.0	14.2	52.7	30.9	72.4	21.6	43.8
Our DDT	76.7	67.1	57.9	30.5	13.0	81.9	48.3	75.7	18.4	48.8	27.5	71.8	66.8	73.7	6.1	18.5	38.0	54.7	78.6	34.6	49.4
Our DDT ⁺	77.9	67.7	61.8	33.8	14.2	82.5	53.0	75.2	18.9	53.5	28.3	73.8	68.7	77.5	8.4	17.6	40.8	55.3	78.6	35.0	51.1

Table 4

Comparisons of the CorLoc metric with weakly supervised object localization methods on VOC 2007. Note that, the “✓” in the “Neg.” column indicates that these WSOL methods require access to a negative image set, whereas our DDT does not.

Methods	Neg.	aero	bike	bird	boat	bottle	bus	car	cat	chair	cow	table*	dog	horse	mbike	person	plant	sheep	sofa	train	tv	Mean
Siva and Xiang [45]	✓	42.4	46.5	18.2	8.8	2.9	40.9	73.2	44.8	5.4	30.5	19.0	34.0	48.8	65.3	8.2	9.4	16.7	32.3	54.8	5.5	30.4
Shi et al. [46]	✓	67.3	54.4	34.3	17.8	1.3	46.6	60.7	68.9	2.5	32.4	16.2	58.9	51.5	64.6	18.2	3.1	20.9	34.7	63.4	5.9	36.2
Cimbis et al. [47]	✓	56.6	58.3	28.4	20.7	6.8	54.9	69.1	20.8	9.2	50.5	10.2	29.0	58.0	64.9	36.7	18.7	56.5	13.2	54.9	59.4	38.8
Wang et al. [15]	✓	37.7	58.8	39.0	4.7	4.0	48.4	70.0	63.7	9.0	54.2	33.3	37.4	61.6	57.6	30.1	31.7	32.4	52.8	49.0	27.8	40.2
Bilen et al. [23]	✓	66.4	59.3	42.7	20.4	21.3	63.4	74.3	59.6	21.1	58.2	14.0	38.5	49.5	60.0	19.8	39.2	41.7	30.1	50.2	44.1	43.7
Ren et al. [48]	✓	79.2	56.9	46.0	12.2	15.7	58.4	71.4	48.6	7.2	69.9	16.7	47.4	44.2	75.5	41.2	39.6	47.4	32.2	49.8	18.6	43.9
Wang et al. [24]	✓	80.1	63.9	51.5	14.9	21.0	55.7	74.2	43.5	26.2	53.4	16.3	56.7	58.3	69.5	14.1	38.3	58.8	47.2	49.1	60.9	47.7
Bilen and Vedaldi [43]	✓	68.9	68.7	65.2	42.5	40.6	72.6	75.2	53.7	29.7	68.1	33.5	45.6	65.9	86.1	27.5	44.9	76.0	62.4	66.3	66.8	58.0
Teh et al. [44]	✓	84.0	64.6	70.0	62.4	25.8	80.6	73.9	71.5	35.7	81.6	46.5	71.3	79.1	78.8	56.7	34.3	69.8	56.7	77.0	72.7	64.6
Our DDT		67.3	63.3	61.3	22.7	8.5	64.8	57.0	80.5	9.4	49.0	22.5	72.6	73.8	69.0	7.2	15.0	35.3	54.7	75.0	29.4	46.9
Our DDT+		71.4	65.6	64.6	25.5	8.5	64.8	61.3	80.5	10.3	49.0	26.5	72.6	75.2	69.0	9.9	12.2	39.7	55.7	75.0	32.5	48.5

Table 5

Comparisons of the CorLoc metric with state-of-the-arts on image sets disjoint with ImageNet.

Methods	chipm.	rhino	stoat	raccoon	rake	wheelc.	Mean
Cho et al. [14]	26.6	81.8	44.2	30.1	8.3	35.3	37.7
SCDA [8]	32.3	71.6	52.9	34.0	7.6	28.3	37.8
Li et al. [6]	44.9	81.8	67.3	41.8	14.5	39.3	48.3
Our DDT	70.3	93.2	80.8	71.8	30.3	68.2	69.1
Our DDT+	72.8	93.2	80.8	75.7	28.3	71.7	70.4

Table 6

Comparisons of webly-supervised classification on Web-Cars [30].

Methods	Strategy	Accuracy
Simple-CNN	GAP	66.86
Zhuang et al. [30]	Attention	76.58
Ours (thr=0)	DDT → GAP	69.79
Ours (thr=0)	DDT → Attention	76.18
Ours (thr=0.1)	DDT → GAP	71.66
Ours (thr=0.1)	DDT → Attention	78.92

Table 7

Comparisons of webly-supervised classification on WebImageNet [30].

Methods	Strategy	Accuracy
Simple-CNN	GAP	58.81
Zhuang et al. [30]	Attention+Neg ^a	71.24
Ours (thr=0)	DDT → GAP	62.31
Ours (thr=0)	DDT → Attention	69.50
Ours (thr=0.1)	DDT → GAP	65.59
Ours (thr=0.1)	DDT → Attention	73.06

^a In the experiments on WebImageNet of [30], beyond attention, they also incorporated 5000 negative class web images for reducing noise. However, we do not require any negative images.

subsets) disjoint with the images from ImageNet. Note that, the six categories (i.e., “chipmunk”, “rhino”, “stoat”, “raccoon”, “rake” and “wheelchair”) of these images are unseen by pre-trained models. The six subsets were provided in [6]. Table 5 presents the CorLoc metric on these subsets. Our DDT (69.1% on average) and DDT+ (70.4% on average) still significantly outperform other methods on all categories, especially for some difficult objects categories, e.g., *rake* and *wheelchair*. In addition, the mean CorLoc metric of [6] based on VGG-19 is only 51.6% on this dataset.

Furthermore, in Fig. 4, several successful predictions by DDT and also some failure cases on this dataset are provided. In particular, for “rake” (“wheelchair”), even though a large portion of images in these two categories contain both people and rakes (wheelchairs), our DDT could still accurately locate the common object in all the images, i.e., rakes (wheelchairs), and ignore people. This observation validates the effectiveness (especially for the high CorLoc metric on *rake* and *wheelchair*) of our method from the qualitative perspective.

4.5. Detecting noisy images

In this section, we quantitatively present the ability of the proposed DDT method to identify noisy images. As aforementioned, in *Object Discovery*, there are 18%, 11% and 7% noisy images in the corresponding categories. In our DDT, the number of positive values in P^1 can be interpreted as a detection score. The lower the number is, the higher the probability of noisy images will be. In particular, no positive value at all in P^1 presents the image as definitely a noisy image. For each category in that dataset, the ROC curve is shown in Fig. 5, which measures how the methods correctly detect noisy images. In the literature, only the method

Table 8

Comparisons of detection results on the VOC 2007 test set. Note that, “07+12” presents the training data is the union set of VOC 2007 trainval and VOC 2012 trainval. “COCO” denotes that the COCO trainval set is used for training. “DDT” denotes that the webly data processed by DDT augmentation is used for training.

Data	aero	bike	bird	boat	bottle	bus	car	cat	chair	cow	table*	dog	horse	mbike	person	plant	sheep	sofa	train	tv	mAP (%)
07+12	76.5	79.0	70.9	65.5	52.1	83.1	84.7	86.4	52.0	81.9	65.7	84.8	84.6	77.5	76.7	38.8	73.6	73.9	83.0	72.6	73.2
COCO ^a +07+12	84.3	82.0	77.7	68.9	65.7	88.1	88.4	88.9	63.6	86.3	70.8	85.9	87.6	80.1	82.3	53.6	80.4	75.8	86.6	78.9	78.8
DDT+07+12	77.6	82.2	77.2	64.9	61.2	85.4	87.2	88.6	58.2	82.6	69.7	85.9	87.0	78.9	78.5	46.3	76.6	73.5	82.5	75.1	76.0

^a Note that, the COCO trainval set contains **120k human labeled** images involving 80 object categories. While, our DDT augmentation only depends on **10k** images of 20 object categories, in especial, these images are **automatically labeled** by the proposed DDT method.

Table 9

Comparisons of detection results on the VOC 2012 test set. Note that, “07++12” presents the training data is the union set of VOC 2007 trainval+test and VOC 2012 trainval. “COCO” denotes that the COCO trainval set is used for training. “DDT” denotes that the webly data processed by DDT augmentation is used for training.

Data	aero	bike	bird	boat	bottle	bus	car	cat	chair	cow	table*	dog	horse	mbike	person	plant	sheep	sofa	train	tv	mAP (%)
07++12	84.9	79.8	74.3	53.9	49.8	77.5	75.9	88.5	45.6	77.1	55.3	86.9	81.7	80.9	79.6	40.1	72.6	60.9	81.2	61.5	70.4
COCO ^a +07++12	87.4	83.6	76.8	62.9	59.6	81.9	82.0	91.3	54.9	82.6	59.0	89.0	85.5	84.7	84.1	52.2	78.9	65.5	85.4	70.2	75.9
DDT+07++12	86.5	81.9	76.2	63.4	55.4	80.8	80.1	89.7	51.6	78.6	56.2	88.8	84.8	85.5	82.6	50.6	78.1	64.1	85.6	68.1	74.4

^a Note that, the COCO trainval set contains **120k human labeled** images involving 80 object categories. While, our DDT augmentation only depends on **10k** images of 20 object categories, in especial, these images are **automatically labeled** by the proposed DDT method.

in [20] (i.e., the Image-Box model in that paper) could solve object co-localization with noisy data. From these figures, it is apparent to see that, in object co-localization, our DDT has significantly better performance in detecting noisy images than Image-Box (whose noisy detection results are obtained by re-running the publicly available code released by the authors). Meanwhile, our mean CorLoc metric without noise is about 12% higher than theirs on *Object Discovery*, cf. Table 1.

4.6. DDT Augmentation based on web images

As validated by previous experiments, DDT can accurately detect noisy images and meanwhile supply object bounding boxes of images (except for noisy images). Therefore, we can use DDT to process web images. In this section, we report the results of both image classification and object detection when using DDT as a tool for generating valid external data sources from free but noisy web data. This DDT based strategy is denoted as DDT augmentation.

4.6.1. Webly-supervised classification

For web based image classification, we compare DDT augmentation with the current state-of-the-art webly-supervised classification method proposed by [30]. As discussed in the related work, [30] proposed a group attention framework for handling web data. In their method, it employed two level attentions: the first level is designed as the group attention for filtering out noise, and the second level attention is based on the single image for capturing discriminative regions of each image.

In the experiments, we test the methods on the *WebCars* and *WebImageNet* datasets which are also proposed by [30]. In *WebCars*, there are 213,072 car images of totally 431 car model categories collected from web. In *WebImageNet*, [30] used 100 sub-categories of the original *ImageNet* as the categories of their *WebImageNet* dataset. There are 61,639 images belonging to the 100 sub-categories from web in total.

In our DDT augmentation, as what we do in Section 4.5, we first use DDT to obtain the number of positive values in P^1 as the detection score for each image in every category. Here, we divide the detection score by the total number of values in P^1 as the noise rate which is in the range of [0, 1]. The more the noise rate is close to zero, the higher the probability of noisy images will be. In the following, we conduct experiments with two thresholds (i.e., 0 or 0.1) with respect to the noise rate. If the noise rate of an image equals to or is smaller than the threshold, that image will be regarded as a noisy image. Then, we remove it from the original webly dataset. After doing the above processing for every category, we can obtain a relatively clean training dataset. Finally, we train deep CNN networks on that clean dataset. The other specific experimental settings of these two webly datasets follow [30].

Two kinds of deep CNN networks are conducted as the test bed for evaluating the classification performance on both two webly datasets:

- “GAP” represents the CNN model with Global Average Pooling as its last layer before the classification layer (i.e., fc+sigmoid), which is commonly used for the image classification task, e.g., [49] and [50].
- “Attention” represents the CNN model with the attention mechanism on the single image level. Because the method proposed in [30] is equipped with the single image attention strategy, we also compare our method based on this baseline model for fair comparisons.

The quantitative comparisons of our DDT augmentation with [30] are shown in Table 6 and Table 7. In these tables, for example, “DDT \rightarrow GAP” denotes that we first deploy DDT augmentation

and then use the GAP model to conduct classification. As shown in these two tables, for both two base models (i.e., “GAP” and “Attention”), our DDT augmentation with 0.1 threshold performs better than DDT augmentation with 0 threshold, which is reasonable. Because in many cases, the noisy images still contains several related concept regions, these (small) regions might be detected as a part of common objects. Therefore, if we set the threshold as 0.1, this kind of noisy images will be omitted. It will bring more satisfactory classification accuracy. Several detected noisy images by DDT of *WebCars* are listed in Fig. 6.

Comparing with the state-of-the-art (i.e., [30]), our DDT augmentation with 0.1 threshold outperforms it and the GAP baseline apparently, which validate the generalization ability and the effectiveness of the proposed DDT in real-life computer vision tasks, i.e., DDT augmentation in webly-supervised classification. Meanwhile, our DDT method is easy to implement and has low computational cost, which ensures its scalability and usability in the real-world scenarios.

4.6.2. Webly-supervised detection

For web based object detection, we first collect an external dataset from the Internet by Google image search engine, named *WebVOC*, using the categories of the *PASCAL VOC* dataset [9]. In total, we collect 12,776 noisy web images, which has a similar scale as the original *PASCAL VOC* dataset. As the results shown in webly-supervised classification, DDT with 0.1 threshold could be the optimal option for webly noisy images. Firstly, we also use DDT with 0.1 threshold to remove the noisy images for the images belonging to 20 categories in *WebVOC*. Then, 10,081 images are remaining as valid images. Furthermore, DDT are used to automatically generate the corresponding object bounding box for each image. The generated bounding boxes by DDT are regarded as the object “ground truth” bounding boxes for our *WebVOC* detection dataset. Several random samples of our *WebVOC* dataset with the corresponding DDT generating bounding boxes are shown in Fig. 7.

After that, a state-of-the-art object detection method, i.e., Faster R-CNN [51], is trained as the base model on different training data to validate the effectiveness of DDT augmentation on the object detection task. For the test sets of detection, we employ the *VOC 2007* and *VOC 2012* test set and report the results in Table 8 and Table 9, respectively.

For testing on *VOC 2007*, following [51], Faster R-CNN is trained on “07+12” and “COCO+07+12”. “07+12” presents the training data is the union set of *VOC 2007* trainval and *VOC 2012* trainval. “COCO+07+12” denotes that except for *VOC 2007* and *VOC 2012*, the *COCO* trainval set is also used for training. “DDT+07+12” is our proposal, which uses DDT to process the web images as aforementioned and then combines the processed web data with “07+12” as the final training data.

As shown in Table 8, our proposal outperforms “07+12” by 2.8% on *VOC 2007*, which is a large margin on the object detection task. In addition, the detection mAP of DDT augmentation is 4% better than “07+12” on the *VOC 2012* test set, cf. Table 9. Note that, our DDT augmentation only depends on **10k** images of 20 object categories, in especial, these images are **automatically labeled** by the proposed DDT method.

On the other hand, our mAP is comparable with the mAP training on “COCO+07+12” in Table 8 (or “COCO+07+12” in Table 9). Here, we would like to point out that the *COCO* trainval set contains **120k human labeled** images involving 80 object categories, which requires much more human labors, capital and time costs than our DDT augmentation. Therefore, these detection results could validate the effectiveness of DDT augmentation on the object detection task.

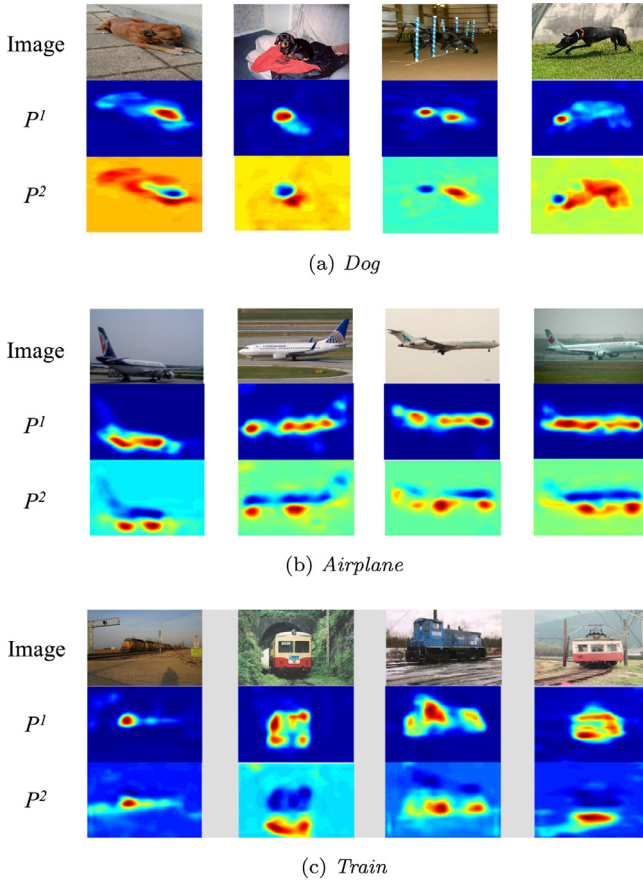


Fig. 8. Four images belonging to each of three categories of VOC 2007 with visualization of their indicator matrices P^1 and P^2 . In visualization figures, warm colors indicate positive values, and cool colors present negative. (Best viewed in color.).

4.7. Further study

In the above, DDT only utilizes the information of the first principal components, i.e., P^1 . How about others, e.g., the second principal components P^2 ? In Fig. 8, we show four images from each of three categories (i.e., dogs, airplanes and trains) in PASCAL VOC with the visualization of their P^1 and P^2 . Through these figures, it is apparently to find P^1 can locate the whole common object. However, P^2 interestingly separates a part region from the main object region, e.g., the head region from the torso region for *dogs*, the wheel and engine regions from the fuselage region for *airplanes*, and the wheel region from the train body region for *trains*. Meanwhile, these two meaningful regions can be easily distinguished from the background. These observations inspire us to use DDT for the more challenging *part-based* object co-localization task in the future, which is never touched before in the literature.

5. Conclusions

Pre-trained models are widely used in diverse applications in pattern recognition and computer vision. However, the treasures beneath pre-trained models are not exploited sufficiently. In this paper, we proposed Deep Descriptor Transformation (DDT) for object co-localization. DDT indeed revealed another reusability of deep pre-trained networks, i.e., convolutional activations/descriptors can play a role as a common object detector. It offered further understanding and insights about CNNs. Besides, our proposed DDT method is easy to implement, and it achieved great object co-localization performance. Moreover, the general-

ization ability and robustness of DDT ensure its effectiveness and powerful reusability in real-world applications. Thus, DDT can be used to handle free but noisy web images and further generate valid data sources for improving both recognition and detection accuracy.

Additionally, DDT also has the potential ability in the applications of video-based unsupervised object discovery. Meanwhile, by considering the fruitful webly video data sources, DDT could be one of the most important parts of a solution for life-long learning. Furthermore, interesting observations in Section 4.7 make the more challenging but intriguing part-based object co-localization problem be a future work.

Acknowledgments

This research was supported by the program A for Outstanding PhD candidate of Nanjing University (201702A010), the National Natural Science Foundation of China under Grant 61772256, Grant 61422203 and Grant 61751306. The authors would like to thank the anonymous reviewers, whose comments have helped improving this paper.

References

- [1] Z.-H. Zhou, Learnware: on the future of machine learning, *Front. Comp. Sci.* 10 (4) (2016) 589–590.
- [2] O. Russakovsky, J. Deng, H. Su, J. Krause, S. Satheesh, S. Ma, Z. Huang, A. Karpathy, A. Khosla, M. Bernstein, A.C. Berg, L. Fei-Fei, Imagenet large scale visual recognition challenge, *IJCV* 115 (3) (2015) 211–252.
- [3] S.M. Erfani, S. Rajasegarar, S. Karunasekera, C. Leckie, High-dimensional and large-scale anomaly detection using a linear one-class SVM with deep learning, *Pattern Recognit.* 58 (2016) 121–134.
- [4] S.H. Lee, C.S. Chan, S.J. Mayo, P. Remagnino, How deep learning extracts and learns leaf features for plant classification, *Pattern Recognit.* 71 (2017) 1–13.
- [5] M. Cimpoi, S. Maji, I. Kokkinos, A. Vedaldi, Deep filter banks for texture recognition, description, and segmentation, *IJCV* 118 (1) (2016) 65–94.
- [6] Y. Li, L. Liu, C. Shen, A.V.D. Hengel, Image co-localization by mimicking a good detector's confidence score distribution, in: B. Leibe, J. Matas, N. Sebe, M. Welling (Eds.), *ECCV, LNCS 9906*, Springer, Cham, 2016, pp. 19–34.
- [7] A. Ghodrati, A. Diba, M. Pedersoli, T. Tuytelaars, L. Van Gool, Deep proposals: hunting objects and actions by cascading deep convolutional layers, *IJCV* 124 (2) (2017) 115–131.
- [8] X.-S. Wei, J.-H. Luo, J. Wu, Z.-H. Zhou, Selective convolutional descriptor aggregation for fine-grained image retrieval, *IEEE TIP* 26 (6) (2017) 2868–2881.
- [9] M. Everingham, S.M.A. Eslami, L.V. Gool, C.K.L. Williams, J. Winn, A. Zisserman, The PASCAL visual object classes challenge: a retrospective, *IJCV* 111 (1) (2015) 98–136.
- [10] X.-S. Wei, C.-L. Zhang, Y. Li, C.-W. Xie, J. Wu, C. Shen, Z.-H. Zhou, Deep descriptor transforming for image co-localization, in: *IJCAI*, 2017, pp. 3048–3054.
- [11] B. Zhou, A. Khosla, A. Lapedriza, A. Oliva, A. Torralba, Learning deep features for discriminative localization, in: *CVPR*, 2016, pp. 2921–2929.
- [12] B. Zhou, A. Khosla, A. Lapedriza, A. Oliva, A. Torralba, Object detectors emerge in deep scene CNNs, in: *ICLR*, 2015, pp. 1–12.
- [13] I. Huerfía, M. Pedersoli, J. Gonzalez, A. Sanfeliu, Combining where and what in change detection for unsupervised foreground learning in surveillance, *Pattern Recognit.* 48 (3) (2015) 709–719.
- [14] M. Cho, S. Kwak, C. Schmid, J. Ponce, Unsupervised object discovery and localization in the wild: Part-based matching with bottom-up region proposals, in: *CVPR*, 2015, pp. 1201–1210.
- [15] X. Wang, Z. Zhu, C. Yao, X. Bai, Relaxed multiple-instance SVM with application to object discovery, in: *ICCV*, 2015, pp. 1224–1232.
- [16] X. Wang, Z. Zhang, Y. Ma, X. Bai, W. Liu, Z. Tu, Robust subspace discovery via relaxed rank minimization, *Neural Comput.* 26 (3) (2014) 611–635.
- [17] H. Zhao, Y. Fu, Semantic single video segmentation with robust graph representation, in: *IJCAI*, 2015, pp. 2219–2225.
- [18] G. Kim, E.P. Xing, L. Fei-Fei, T. Kanade, Distributed co-segmentation via submodular optimization on anisotropic diffusion, in: *ICCV*, 2011, pp. 169–176.
- [19] A. Joulin, F. Bach, J. Ponce, Multi-class co-segmentation, in: *CVPR*, 2012, pp. 139–150.
- [20] K. Tang, A. Joulin, L. Li, L. Fei-Fei, Co-localization in real-world images, in: *CVPR*, 2014, pp. 1464–1471.
- [21] P. Tang, X. Wang, Z. Huang, X. Bai, W. Liu, Deep patch learning for weakly supervised object classification and discovery, *Pattern Recognit.* 71 (2017) 446–459.
- [22] D. Zhang, D. Meng, L. Zhao, J. Han, Bridging saliency detection to weakly supervised object detection based on self-paced curriculum, in: *IJCAI*, 2016, pp. 3538–3544.
- [23] H. Bilen, M. Pedersoli, T. Tuytelaars, Weakly supervised object detection with convex clustering, in: *CVPR*, 2015, pp. 1081–1089.

- [24] C. Wang, W. Ren, K. Huang, T. Tan, Large-scale weakly supervised object localization via latent category learning, *IEEE TIP* 24 (4) (2015) 1371–1385.
- [25] P. Tang, X. Wang, X. Bai, W. Liu, Multiple instance detection network with on-line instance classifier refinement, in: *CVPR*, 2017, pp. 2843–2851.
- [26] A. Joulin, K. Tang, L. Fei-Fei, Efficient image and video co-localization with Frank-Wolfe algorithm, in: D. Fleet, T. Pajdla, B. Schiele, T. Tuytelaars (Eds.), *ECCV, LNCS 8694*, Springer, Cham, 2014, pp. 253–268.
- [27] T.-Y. Lin, M. Maire, S. Belongie, J. Hays, P. Perona, D. Ramanan, P. Dollár, C.L. Zitnick, Microsoft COCO: Common objects in context, in: D. Fleet, T. Pajdla, B. Schiele, T. Tuytelaars (Eds.), *ECCV, LNCS 8694*, Springer, Cham, 2014, pp. 740–755.
- [28] S. Ma, S.A. Bargal, J. Zhang, L. Sigal, S. Sclaroff, Do less and achieve more: training CNNs for action recognition utilizing action images from the web, *Pattern Recognit.* 68 (2017) 334–345.
- [29] K. Yu, Z. Wang, L. Zhuo, J. Wang, Z. Chi, D. Feng, Learning realistic facial expressions from web images, *Pattern Recognit.* 46 (8) (2013) 2144–2155.
- [30] B. Zhuang, L. Liu, Y. Li, C. Shen, I. Reid, Attend in groups: A weakly-supervised deep learning framework for learning from web data, in: *CVPR*, 2017, pp. 1878–1887.
- [31] G. Papandreou, L.-C. Chen, K. Murphy, A.L. Yuille, Weakly- and semi-supervised learning of a DCNN for semantic image segmentation, in: *ICCV*, 2015, pp. 1742–1750.
- [32] T. Xiao, T. Xia, Y. Yang, C. Huang, X. Wang, Learning from massive noisy labeled data for image classification, in: *CVPR*, 2015, pp. 2691–2699.
- [33] R. Fergus, P. Perona, A. Zisserman, A visual category filter for Google images, in: T. Pajdla, J. Matas (Eds.), *ECCV, LNCS 3021*, Springer, Verlag, 2004, pp. 242–256.
- [34] X.-J. Wang, L. Zhang, X. Li, W.-Y. Ma, Annotating images by mining image search results, *IEEE TPAMI* 30 (11) (2008) 1919–1932.
- [35] K. Pearson, On lines and planes of closest fit to systems of points in space, *Philos. Mag.* 2 (11) (1901) 559–572.
- [36] M.D. Zeiler, R. Fergus, Visualizing and understanding convolutional networks, in: D. Fleet, T. Pajdla, B. Schiele, T. Tuytelaars (Eds.), *ECCV, LNCS 8694*, Springer, Cham, 2014, pp. 818–833.
- [37] B. Hariharan, P. Arbeláez, R. Girshick, J. Malik, Object instance segmentation and fine-grained localization using hypercolumns, *IEEE TPAMI* (4) (2017) 627–639.
- [38] E. Shelhamer, J. Long, T. Darrell, Fully convolutional networks for semantic segmentation, *IEEE TPAMI* 39 (4) (2017) 640–651.
- [39] K. Simonyan, A. Zisserman, Very deep convolutional networks for large-scale image recognition, in: *ICLR*, 2015, pp. 1–14.
- [40] A. Vedaldi, K. Lenc, MatConvNet - convolutional neural networks for MATLAB, in: *ACM MM*, 2015, pp. 689–692.
- [41] M. Rubinstein, A. Joulin, J. Kopf, C. Liu, Unsupervised joint object discovery and segmentation in internet images, in: *CVPR*, 2013, pp. 1939–1946.
- [42] A. Joulin, F. Bach, J. Ponce, Discriminative clustering for image co-segmentation, in: *ICCV*, 2010, pp. 2984–2991.
- [43] H. Bilen, A. Vedaldi, Weakly supervised deep detection networks, in: *CVPR*, 2016, pp. 2846–2854.
- [44] E.W. Teh, M. Ročan, Y. Wang, Attention networks for weakly supervised object localization, in: *BMVC*, 2016, pp. 1–11.
- [45] P. Siva, T. Xiang, Weakly supervised object detector learning with model drift detection, in: *CVPR*, 2011, pp. 343–350.
- [46] Z. Shi, T.M. Hospedales, T. Xiang, Bayesian joint modelling for object localisation in weakly labelled images, *IEEE TPAMI* 37 (10) (2015) 1959–1972.
- [47] R.G. Cinbis, J.J. Verbeek, C. Schmid, Weakly supervised object localization with multi-fold multiple instance learning, *IEEE TPAMI* 39 (1) (2017) 189–203.
- [48] W. Ren, K. Huang, D. Tao, T. Tan, Weakly supervised large scale object localization with multiple instance learning and bag splitting, *IEEE TPAMI* 38 (2) (2016) 405–416.
- [49] C. Szegedy, W. Liu, Y. Jia, P. Sermanet, S. Reed, D. Anguelov, D. Erhan, V. Vanhoucke, A. Rabinovich, Going deeper with convolutions, in: *CVPR*, 2015, pp. 1–9.
- [50] K. He, X. Zhang, S. Ren, J. Sun, Deep residual learning for image recognition, in: *CVPR*, 2016, pp. 770–778.
- [51] S. Ren, K. He, R. Girshick, J. Sun, Faster R-CNN: towards real-time object detection with region proposal networks, *IEEE TPAMI* 39 (6) (2017) 1137–1149.

Xiu-Shen Wei received the BSc and PhD degrees in computer science in 2012 and 2018, respectively. He joined the Megvii (Face++) Technology Ltd. in 2017, and is currently the research lead in Megvii Research Nanjing. He achieved the first place in the Apparent Personality Analysis competition (in association with ECCV 2016) and the first runner-up in the Cultural Event Recognition competition (in association with ICCV 2015) as the team director. His research interests are computer vision and machine learning.

Chen-Lin Zhang received his BS degree in the Department of Computer Science and Technology from Nanjing University, China, in 2016. He is currently working toward the PhD degree in the Department of Computer Science and Technology, Nanjing University, China. His research interests are computer vision and machine learning.

Jianxin Wu received his BS and MS degrees in computer science from Nanjing University, and his PhD degree in computer science from the Georgia Institute of Technology. He is currently a professor in the Department of Computer Science and Technology at Nanjing University, China, and is associated with the National Key Laboratory for Novel Software Technology, China. He was an assistant professor in the Nanyang Technological University, Singapore, and has served as an area chair for CVPR 2017, ICCV 2015, senior PC member for AAAI 2017, AAAI 2016 and an associate editor for Pattern Recognition Journal. His research interests are computer vision and machine learning. He is a member of the IEEE.

Chunhua Shen received the degrees from Nanjing University, Australian National University, and the Ph.D. degree from The University of Adelaide. He was with the Computer Vision Program, NICTA (National ICT Australia), Canberra Research Laboratory, for about six years. He is currently a Professor of Computer Science with The University of Adelaide. His research interests are in the intersection of computer vision and statistical machine learning. He received the Australian Research Council Future Fellowship in 2012.

Zhi-Hua Zhou received the BSc, MSc and PhD degrees in computer science from Nanjing University, China, in 1996, 1998 and 2000, respectively, all with the highest honors. He joined the Department of Computer Science & Technology at Nanjing University as an Assistant Professor in 2001, and is currently Chair Professor and Standing Deputy Director of the National Key Laboratory for Novel Software Technology; he is also the Founding Director of the LAMDA group. His research interests are mainly in artificial intelligence, machine learning, data mining and pattern recognition. He is a Fellow of the ACM, AAAS, AAAI, IEEE, IAPR, IET/IEE and CCF.