



# A deep one-shot network for query-based logo retrieval

Ayan Kumar Bhunia<sup>a</sup>, Ankan Kumar Bhunia<sup>b</sup>, Shuvojit Ghose<sup>c</sup>, Abhirup Das<sup>c</sup>,  
Partha Pratim Roy<sup>d,\*</sup>, Umapada Pal<sup>e</sup>

<sup>a</sup> University of Surrey, England, United Kingdom

<sup>b</sup> Jadavpur University, India

<sup>c</sup> Institute of Engineering & Management, India

<sup>d</sup> Indian Institute of Technology Roorkee, India

<sup>e</sup> Indian Statistical Institute, India

## ARTICLE INFO

### Article history:

Received 12 February 2019

Revised 1 July 2019

Accepted 10 July 2019

Available online 11 July 2019

### Keywords:

Logo retrieval

One-shot learning

Multi-scale conditioning

Similarity matching

Query retrieval

## ABSTRACT

Logo detection in real-world scene images is an important problem with applications in advertisement and marketing. Existing general-purpose object detection methods require large training data with annotations for every logo class. These methods do not satisfy the incremental demand of logo classes necessary for practical deployment since it is practically impossible to have such annotated data for new unseen logo. In this work, we develop an easy-to-implement query-based logo detection and localization system by employing a one-shot learning technique using off the shelf neural network components. Given an image of a query logo, our model searches for logo within a given target image and predicts the possible location of the logo by estimating a binary segmentation mask. The proposed model consists of a conditional branch and a segmentation branch. The former gives a conditional latent representation of the given query logo which is combined with feature maps of the segmentation branch at multiple scales in order to obtain the matching location of the query logo in a target image. Feature matching between the latent query representation and multi-scale feature maps of segmentation branch using simple concatenation operation followed by  $1 \times 1$  convolution layer makes our model scale-invariant. Despite its simplicity, our query-based logo retrieval framework achieved superior performance in FlickrLogos-32 and TopLogos-10 dataset over different existing baseline methods.

© 2019 Published by Elsevier Ltd.

## 1. Introduction

Detection of logos in scene images and videos has a number of useful applications: commercial analysis of brands [1], vehicle-logo detection [2] for intelligent traffic-control systems and even Augmented Reality [3]. A logo is a unique symbol representing any brand or organization that expresses its functionality and distinguishes the brand and its products from others. A merchant can easily assess the presence of his brand in television, social media and e-commerce sites by searching for the company's unique logo. Detection and localization of logos is a crucial step in inspecting market trends, allowing companies to meet the customers' needs by optimizing existing marketing schemes. Automating logo detection will also allow merchants to detect copyright infringement cases by testing the originality of suspicious advertisements.

Logo detection problem can be seen as a special case of object detection [4] in images. The appearance of a logo varies drastically

in real-world images due to lighting effects, occlusions, rotations, shearing effects and scales. It is hard to detect different sizes of logos in a diverse contextual environment with uncontrolled illumination, low resolution, and high background clutter. In the recent years, logo detection has gained a lot of attention from the computer vision community [5–10]. Earlier works of logo detection rely on the bag-of-words approach [11] where SIFT features are quantized into a vocabulary of learned logo patterns in images. Boia et al. [6] used a novel approach based on homographic class graphs to perform both logo localization and recognition. Recently, significant improvement has been achieved by adopting deep-learning techniques in this field [7,12,13]. In the meantime, several object detection algorithms have been introduced, namely, R-CNN [14], Fast R-CNN [15] and Faster R-CNN [4] which have been successfully adapted for the logo detection problem [8], boosting object-recognition performances further. In addition, due to advent of deep learning, there have been significant progress in scene-text detection [16], scene text recognition [17], script identification [18] tasks etc; however, there are very limited recent

\* Corresponding author.

E-mail address: [proy.fcs@iitr.ac.in](mailto:proy.fcs@iitr.ac.in) (P.P. Roy).

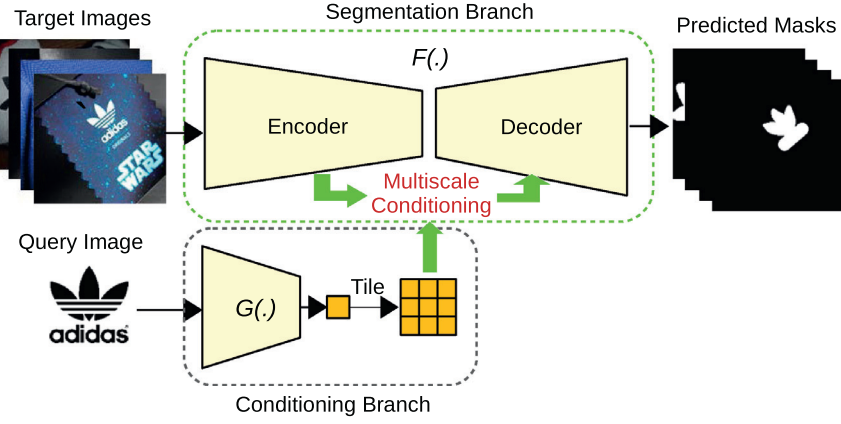


Fig. 1. Illustration of Query-based logo detection problem.

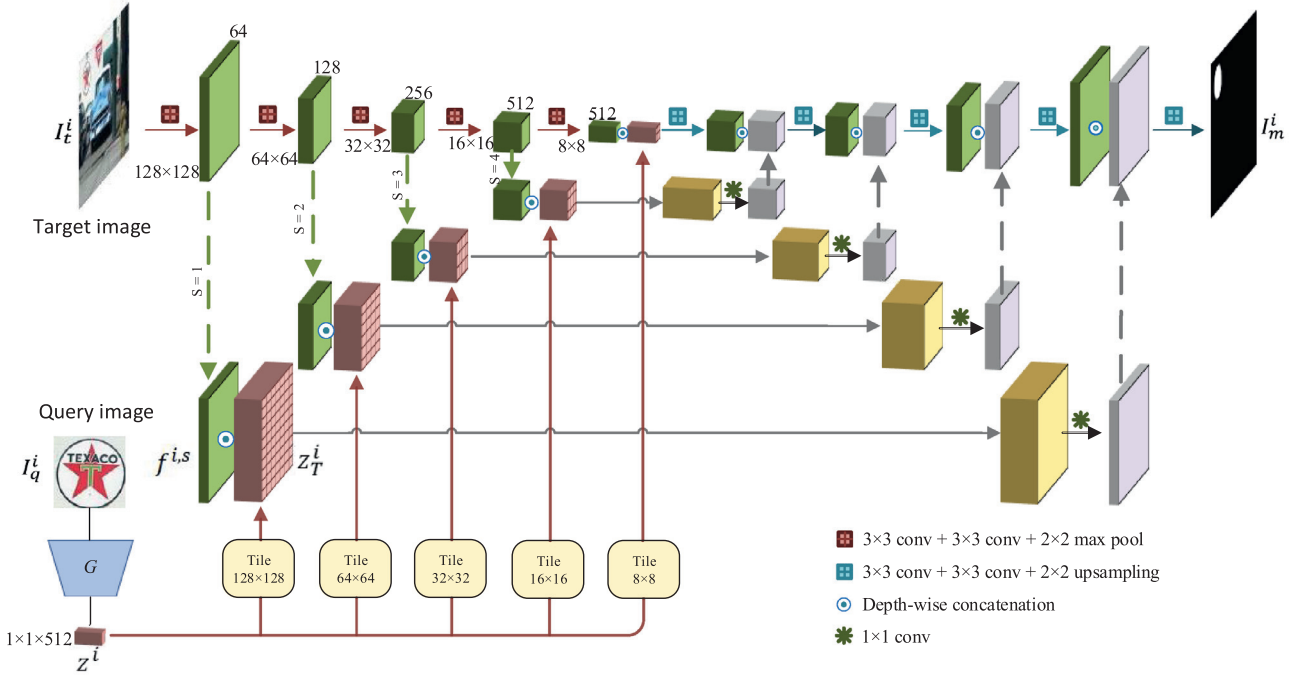


Fig. 2. Query-based logo detection framework: At-first, the conditional network  $G$  takes the query logo image  $I_q^i$  and outputs a latent representation  $z_i^i \in \mathbb{R}^{1 \times 1 \times 512}$ . In the main segmentation branch, the target image  $I_t^i$  is taken as input. After 5 stages of down-sampling, the encoder part reduces the input to  $8 \times 8 \times 256$ . The earlier obtained  $z_i^i$  is concatenated with last layer of each stage via a tile operation. Suppose, at stage-1, i.e.  $s=1$ , the feature map  $f^{1,s}$  is combined with the  $z_i^i$  after the necessary tile operation. Then, after applying a  $1 \times 1$  convolution on the combined representation, it is further concatenated with its respective layer at the decoder side as shown in the figure.

deep-frameworks towards logo detection and recognition in spite of its complexity.

Deep-learning based frameworks are largely data-driven, contrary to logo-datasets that have several image classes but few images. Since deployable logo detection models need to be robust to new unseen logos, the model should be designed to satisfy the incremental demands for logo classes, contrary to existing methods which are limited to a set of seen logos and are not scalable enough for practical deployment. With the current problem setup, it is impossible for models to work with the logo of any brand. To meet the need for a scalable solution, we re-design the problem statement as shown in Fig. 1. Given a query logo of a particular company or organization, the main objective would be to find out whether the logo of the same company or organization is present in a target image or not. If present, the model fetches its position within the given scene.

Following this problem statement, we propose a one-shot learning based technique to design a framework that adapts to

new logo classes in a data-efficient way. Recently, one-shot learning [19] has gained notable attention in learning new concepts from sparse data. One-shot image classification [20] and one-shot image segmentation [21–24] are some practical applications of it. Inspired by these works, we propose a one-shot query based logo detection framework that consists of two modules: the conditioning module and the segmentation module. The conditioning part takes a query logo image and obtains a latent representation which will be used as a conditional input to the segmentation branch. The segmentation network is a basic encoder-decoder architecture. In order to capture the multi-scale correlation between the query image and different regions of the target image, we concatenate the conditional latent representation with multiple layers of the encoder part(segmentation network), followed by  $1 \times 1$  convolution. The obtained representation is further combined with the respective layers of the decoder part in order to better guide the decoder part of segmentation network to generate binary segmentation mask, forming a skip-connection like

architecture [25]. When the model detects high similarity between a query and some particular region of the target image, the network tries to produce high response at the corresponding portion of the segmentation map. To apply our model for a new logo-class, unlike the fine-tuning approach [21] which may require large number of training samples with corresponding ground truth data, our approach needs only one sample of the logo of a concerned company or organisation. This query logo sample can be obtained very easily from the official logo design image or just by cropping an appropriate logo portion from a scene logo image. In this paper, we make following novel contributions:

- We propose a scalable solution for the logo detection problem by re-designing the traditional problem setting. We present a query-based logo search and detection system by employing a simple, fully differentiable one-shot learning framework which can be used for new logo classes without further training the whole network. To the best of our knowledge, ours is the first work to address one-shot query based deep learning framework for logo retrieval which is novel in the literature of logo-retrieval research.
- To deal with the logos of varying sizes, we propose a novel one-shot framework through *multi-scale conditioning* that is specially designed to learn the similarity between the query image and target image at multiple scales and resolutions.

## 2. Related works

*Traditional methods:* Logo-related research has been carried out for over two decades in the area of computer vision and pattern recognition. Earlier, logo recognition and detection were categorized as a specific problem of object detection. The primary models were developed on geometric object features [26] which relied on properties of objects such as lines, vertices, curves and shapes. Later on, properties of pixel value (luminance or color) were introduced as photometric object [27–29] features, which eventually replaced the previous ones. These features were computed locally, could solve the problem of occlusion to an extent and were able to distinguish similar objects better way [27]. Thereafter many logo-related works were carried out with the methods of content-based indexing and retrieval in trademark databases. The main goal is to assist in trademark infringement detection by checking a newly designed trademark with registered logos in archives [30–33]. The task of trademark recognition in videos is inherently harder due to loss of quality of original logos during processing (e.g. color sub-sampling, video interlacing, motion blur, etc.). However in this case, it is assumed that the acquired images are of good quality and moderate distortion free. Kovar et al. [34] applied a heuristic technique to discard sparse or small populated edge regions of the images and analysed the set of significant edges during logo detection. The work in [35] deals with the logos that appear on the rigid planar surfaces having homogeneous colored background in images using Hough Transformation. Color histogram back projection is applied on candidate logo regions [36] to recognize candidate logos. Multidimensional receptive field histograms are also used to perform the task of logo recognition. Here, the most likely logo region is computed for every candidate region. Therefore, if a region in the image does not contain a logo, the identification precision gets reduced.

The traditional logo recognition models are well established on key-point based detectors and descriptors (specially SIFT). SIFT-based models [29] basically take an image and transform it into a large collection of feature vectors that are invariant to affine transformations and even robust to different lighting conditions. One of its main characteristics is the ability to detect stable salient points in the image across multiple scales. On the

basis of this, Lamberto et al. [37] proposed a representation of trademarks and video frame contents using SIFT feature-points mainly targeting to detect, localize and retrieve trademarks in a robust manner, irrespective of irregularities. The classification of retrieved trademarks is analyzed by matching a set of SIFT feature descriptors for each trademark instance with the features detected in every frame of the video. Kleban et al. [38] proposed a logo detection model by clustering matching spatial configurations of frequent local features and introducing spatial pyramid mining. Alexis and Olivier presented a new content-based retrieval framework [39] using a thresholding strategy in order to improve the accuracy of query images. In [40], the author described the logo detection problem as a small object detection problem and solved it by interactive visual object search through mutual information maximization. Romberg and Lienhart [41] exploited a large scale recognition approach using feature bundling. As feature bundles carry more information about the image content than single visual words, they aggregated individual local features into bundles and then detected logos by querying the database of reference images based on features. Based on the analysis of the local features and basic structure, such as edges, curves, triangles, etc. Romberg et al. [42] presented a system by encoding and indexing the spatial layout of local features found in logos. Revaud et al. [43] introduced dedicated correlation-based burstiness model using a down-weight technique for noisy logo detections. Boia et al. [6,44] smartly exploited homographic class graphs to analyze logo localization and recognition tasks. It is noteworthy that they used inverted secondary models to control inverted colors instances. Marcal et al. [45] presents a robust queried-by-example logo retrieval system where logos are compactly described as a variant of the shape context descriptor. They perform k-NN search in the locality sensitive hashing database to retrieve logos. Jianlong et al. [46] exploits local features to form a visual codebook and build an inverted file to accelerate the indexing process. Then several groups are proposed according to the local feature type namely point type, shape type and patch type. Finally, adaptive feature selection with weight updating mechanism is used to perform logo retrieval. Jinqiao et al. [47] used k-means clustering to develop a visual logo dictionary, and next, latent semantic indexing and analysis is used for logo retrieval. Soysal et al. [48] considered spatial similarity of local patterns by utilizing a descriptor for scene logo retrieval.

*Deep learning-based methods:* Bianco et al. [12] applied an unsupervised segmentation algorithm to produce a number of object proposals that are more likely to contain a logo object. Thereafter these object proposals are processed through a query expansion step in order to deal with the variation of logo instances. Finally, pre-trained CNN model with SVM classifier was used for logo recognition. Based on the previous Deep-CNN based pipeline, Eggert et al. [7] trained an SVM classifier on synthetic training examples to compare with a trained classifier on real images. The main objective of this approach was to demonstrate the benefit of using synthetic images during training. Iandola et al. [13] investigated several variations of GoogLeNet architecture including GoogLeNet with global average pooling and auxiliary classifier after each inception layers. Oliveira et al. [8] exploited Fast R-CNN to detect graphic logos. To tackle the limitation of large-scale graphic logo datasets, transfer learning was used to leverage image representation.

Very recently, Bianco et al. [5] proposed a deep learning pipeline investigating the benefits of different training choices such as class-balancing, data augmentation, contrast normalization, sample-weighting and explicit modelling of background class. Su et al. [9] has introduced a framework to generate new training data for logo detection by synthesizing context and thus it intends to increase the robustness against unseen cluttered background. Based on the similar hypothesis, a more advanced framework

[49] using Generative Adversarial Networks has been designed to generate context consistent logo images for training. Su et al. [50] proposed an incremental learning approach which discovers informative training images from noisy web data in order to improve the performance. In contrast to these recent works, we here intend to explore the existing benchmark logo-datasets and design a new one-shot deep framework for query-based logo retrieval.

*One shot learning:* Creditably, one-shot learning requires only a single annotated image to learn a new class. For few-shot learning, only a few examples of a class are needed to generalize knowledge for recognition. In recent years, one-shot learning is applied to various fields of computer vision such as image classification and visual question answering. The Siamese network architecture by Koch et al. [20] has shown that few-shot image classification can outperform several classification baselines for a binary verification task. The siamese network makes use of two shared network to extract features from two input images and a similarity score between the two feature representations decides the correspondence between the input images. Another important recent work in few-shot classification is Matching networks [51] that learns to determine the correct class label for a given query image from unseen categories. Discriminative methods described in [52,53] have the ability to update parameters of a base classifier that has learnt from training classes, while adapting to new classes for the specific task. But the complication in adapting classifiers in this manner is that they are prone to overfitting. Bertinetto et al. [52] trained two-branch networks, where one branch receives an example and generates a set of dynamic parameters and second branch classifies the query image based on those parameters and a set of learnt static parameters. Noh et al. [54] used a similar approach for question answering. Most of the existing works on one-shot learning focus on classification, not structured output. In [21], a simple approach is proposed by authors to perform one-shot semantic segmentation by fine-tuning a pre-trained segmentation network on a labelled image. But, this method is prone to overfitting. Later, Shaban et al. [23] introduced a two-branched network to support dense semantic image segmentation in the one-shot setting. A N-way (classes) few shot segmentation framework based on metric learning and prototype learning has been introduced by Dong et al. [55]. Very recently, Zhang et al. [24] proposed a one-shot semantic segmentation framework with a new mask average pooling operation [24]. These earlier frameworks do not consider multi-scale information, and the performance is limited in real-world logo datasets.

In contrast, instead of overly complicated network design, we use off-the-shelf neural network components in our architecture design that is trainable end-to-end. The contribution of the work lies in following aspects in terms of network design choices: (1) We use a parameterized  $1 \times 1$  convolution layer in order to measure the similarity in high-dimensional space. (2) A multi-scale conditioning operation is used in order to handle the logos of different scale. (3) Skip-connection between encoder-decoder parts of the network is found to be helpful for better information passing.

### 3. Problem setup

The main objective of the logo detection problem is to find out whether a query logo of a particular company or organization is present in a target image or not. If it is present, we obtain a segmentation map containing information about the spatial location of the query in that target image. More formally, given a query image  $I_q$  and a target image  $I_t$  our job is to find out a segmentation map  $I_m$  which is a binary 2D-matrix. The value 1 in the segmentation map represents the region containing the query logo whereas 0 represents background. During training, we have access to a large number of query target-mask triplets  $\{I_q^i, I_t^i, I_m^i\}_{i=1}^N$  where  $I_m^i$  is the semantic segmentation map of  $I_t^i$  given the conditional

query image  $I_q^i$  of various logo classes. In the experiments, we have seen that our model is also able to generalize to unknown logos.

## 4. Proposed framework

### 4.1. Overview

Our proposed end-to-end network can be divided into two steps: (i) conditioning step and (ii) segmentation step. In the conditioning part, we extract a conditional latent representation of the query image through a Convolutional Neural Network or CNN. The segmentation part of our model is a modified version of the U-Net [25] architecture. The encoder intends to extract more robust feature information from the target image and subsequently, the decoder tries to obtain the corresponding segmentation map conditioned on the latent representation of the query logo-image. By the word *conditioning*, we here intend to segment regions of the target image conditioned on the query latent vector; in other words, the regions in target image which have a close similarity with the conditional query logo image. In order to overcome the possibility of drastic mismatch in scale and resolution between the query logo and the target image containing that logo, we encourage the model to learn *multi-scale conditioning* by combining the logo representation at different scales to the encoder part of the U-Net architecture. The proposed one-shot learning framework is fully differentiable and the two-branch architecture can be trained in an end-to-end manner.

### 4.2. Conditioning module

Logo-to-dense feature map conversion is a primary step in our architecture. Examining the power of extracting robust task-specific feature representation of convolutional networks, we use a CNN as the feature extractor. For this purpose, at first, we resize each query logo image to a fixed size of  $64 \times 64$ . Next, we feed the query image  $I_q^i$  to the network  $G(\cdot)$  which encodes the image to a latent representation  $z^i$ . Specially, the logo image is converted into a multichannel feature vector of unit spatial dimension (i.e.  $1 \times 1 \times 512$ ) for further use. Thus,

$$z^i = G(I_q^i; \theta_G), \quad (1)$$

where  $z^i \in \mathbb{R}^{1 \times 1 \times 512}$  and  $\theta_G$  is the parameter of the network. We adapt VGG-16 like architecture consisting of 13 convolution layers with rectified non-linearity (ReLU) activation after each layer. After 6 max-pooling layers, the input image gets converted to  $\frac{1}{2^6}$  shape of feature map giving a final feature representation of dimension  $1 \times 1 \times 512$ .

### 4.3. Multi-scale segmentation module

In this section, we describe our multi-scale segmentation module. Given a target image  $I_t^i$  and the obtained latent representation  $z^i$  of query logo, the segmentation network  $F(\cdot)$  will try to obtain a segmentation map  $\bar{I}_m^i$ . Mathematically,

$$\bar{I}_m^i = F(I_t^i, z^i; \theta_F), \quad (2)$$

where  $\theta_F$  is the parameters of the network. Following the architecture of U-Net, we have introduced two networks: Encoder network and Decoder network. Encoder network encodes the target image to a latent representation. The encoder follows the typical structure of a CNN. It consists of a set of two convolution layers of  $3 \times 3$  filter size and followed by a  $2 \times 2$  max-pooling layer with stride 2. This set of operations are repeated five times to progressively down-sample the input image until we get a bottleneck layer where we get the final latent feature representation.



In this five-stage down-sampling process we have used increasing number of filters: 64, 128, 256, 512 and 512 respectively in each stage. After, each stage of down-sampling the spatial dimension of the feature maps are reduced  $\frac{1}{2}$  times. Thus, the final latent feature map is  $\frac{1}{2^5}$  of the input image. For example, if  $I_t$  has the shape of  $256 \times 256 \times 3$ , then the latent feature map is of  $8 \times 8 \times 512$ . We have seen that further down-sampling the obtained spatial representation eventually degrade the results to some extent.

In order to learn the correlation between the query image and the target image, a naive approach would be to just concatenate the latent representations of the two networks. However, the said method has several disadvantages: 1. The scale difference between the query and the target image makes it difficult to achieve robust segmentation with a single scale conditioning approach. Thus, multi-scale aggregation is required to obtain detailed parsing maps. 2. Even state-of-the-art logo detection frameworks find it hard to detect smaller logos and ones present in very large scenes. Recent works on object detection [56–59] reveal that the feature map of the shallower convolution layers have higher resolution and are helpful to detect small objects whereas the deeper layers contain richer task-specific semantic information.

**Multi-scale conditioning:** Taking cues from the aforesaid observations, we make use of the in-network hierarchical feature produced from feature maps of encoder part of the segmentation network having different spatial resolution. Let, given a target image  $I_t^i \in \mathbb{R}^{H \times W \times C}$ , the feature map extracted from the  $s$ th stage of the five-stage convolution network be  $f^{i,s} \in \mathbb{R}^{H_s \times W_s \times C_s}$ , where,  $H_s = \frac{H}{2^s}$  and  $W_s = \frac{W}{2^s}$ . Earlier, we got  $z^i \in \mathbb{R}^{1 \times 1 \times 512}$  as the conditional query logo representation. In order to impose condition on the segmentation network using latent query vector  $z^i$ , we use a technique similar to sliding window based template matching. To measure the similarity between latent query vector  $z^i$  and each discrete feature  $f_{mn}^{i,s}$  in  $f^{i,s}$ , we use a simple concatenation operation followed by a  $1 \times 1$  convolution.

$$\bar{f}_{mn}^{i,s} = \text{Conv}_{1 \times 1}^s([f_{mn}^{i,s}; z^i]), \quad (3)$$

where  $f_{mn}^{i,s}$  represents the feature representation at position  $(m, n)$ , and which, upon concatenation with  $z^i$  gives a vector of size  $1 \times 1 \times (C_s + 512)$ . The number of  $1 \times 1$  convolutional filters in  $s$ th layer is set to  $C_s$ , and thus it gives  $\bar{f}_{mn}^{i,s}$  of dimension  $1 \times 1 \times C_s$ . By this simple operation (Eq. (3)), it tries to capture the similarity between the latent query vector and feature embedding at  $(m, n)$  position of the feature map  $f^{i,s}$  from segmentation network. High activation values in  $\bar{f}^{i,s}$  signifies high extent of matching that the query logo is likely to be present at that location and vice-versa. Mathematically, let's denote conditional query representation as  $Z$  and one discrete feature  $f_{mn}^{i,s}$  as  $F$ . Using  $1 \times 1$  convolution operation with parameters  $\theta_{1 \times 1}$ , we intend to model  $P(Y|F, Z; \theta_{1 \times 1})$  where  $Y$  denotes the similarity between  $F$  and  $Z$ . To illustrate,  $P(Y|F_1, Z_1; \theta_{1 \times 1}) > P(Y|F_2, Z_2; \theta_{1 \times 1})$  when the similarity between  $F_1$  and  $Z_1$  is higher than  $F_2$  and  $Z_2$ . These parameters of  $1 \times 1$  convolutions are learnt implicitly from the traditional cross-entropy loss in an end-to-end manner. In other words, this  $1 \times 1$  convolution is a function that takes  $Z$  and  $F$  as input, and predicts their similarity at its output in terms of activation values. Unlike other works [24], instead of using cosine-similarity computation for feature matching, we make such choice mainly because of two reasons: : *First*, using a parameterized layer to capture the similarity is expected to be robust to various kind of deformations, illuminations present in real word logo retrieval scenarios. *Second*,  $1 \times 1$  convolution handles the intra-class variance well, in other words, sometimes it is found that despite same logo-class it has some difference in terms of the color of the logo, font used, presence of text inside a logo (for example in *adidas* logos, sometimes the text remains missing). Also, it is to be noted that we use the final latent query representation

from conditioning module for multi-scale conditioning at every layer in encoder part (segmentation network) since deeper latent representation is supposed to contain better semantic information of the query logo, instead of using pooled feature representation from earlier layers of conditioning module. More comparison with alternative choices are given in ablation study (Section 5.5).

For faster and efficient implementation, a tile operation is performed on  $z^i$  to convert it to a specific spatial dimension such that is compatible for concatenation with  $f^{i,s}$ :  $z_T^{i,s} = \text{Tile}(z^i; s_c)$ , where  $s_c$  is the scale required to tile  $z^i \in \mathbb{R}^{1 \times 1 \times 512}$  to a new dimension  $z_T^{i,s} \in \mathbb{R}^{H_s \times W_s \times 512}$  and concatenated with  $f^{i,s}$  giving a dimension of  $\mathbb{R}^{H_s \times W_s \times 512 + C_s}$ . This is followed by a  $1 \times 1$  convolution with  $C_s$  filters:

$$\bar{F}^{i,s} = \text{Conv}_{1 \times 1}^s([f^{i,s}; z_T^{i,s}]) \quad (4)$$

where  $\bar{F}^{i,s} \in \mathbb{R}^{H_s \times W_s \times C_s}$ . Alternatively, it can be interpreted as applying a sliding window based template matching over the target image with stride  $2^s$  (since every max-pooling layers steps down the spatial resolution of feature-map by a factor of 2) for  $s$ th stage in encoder part of segmentation network.

The decoder network of the segmentation model is similar to the encoder part, the only difference being that up-convolution has been used. Every stage in the decoder path consists of an up-sampling layer followed by a  $2 \times 2$  up-convolution layer. After that, we combine this with the previously obtained fused representation of the corresponding stage. At last, two  $3 \times 3$  convolutions, followed by a ReLU completes the set of operations for a particular stage. These set of operations are repeated for all the stages until we get the predicted segmentation map  $\bar{I}_m$  through a SoftMax layer. At the end, we define a pixel-wise binary cross entropy loss  $L$  which is used to train the complete model.

$$L = \frac{1}{HW} \sum_{i=0}^W \sum_{j=0}^H \left( -I_m^i \log(\bar{I}_m^i) \right) \quad (5)$$

The aim of the model is to minimize this loss by updating the parameter of the both the conditioning ( $\theta_C$ ) and segmentation ( $\theta_F$ ) modules of the network through back-propagation technique in an end-to-end manner.

Some earlier works like Mask-RCNN [60] first predicts the bounding box using Region Proposal Network(RPN) and then performs the instance segmentation as a two steps process. This requires two separate loss functions for region proposal and final instance segmentation, respectively. In contrast, we avoid using a two steps process and predict the segmentation map in a single step directly, and thus it makes the process faster as well as easy for implementation. That becomes feasible because of *multiscale conditioning* and *parameterized similarity matching layer* with *skip connections* that helps in better information passing between encoder and decoder part of the framework.

## 5. Experiment

We have evaluated the performance of our framework in both one-shot and traditional settings. In the one-shot setting, given an unseen conditional query logo image (from the test set), the model trained on a completely disjoint set of logo classes returns a reference binary mask for localization. In the traditional setting, on top of our query-based strategy, we experiment with the traditional train-test split strategy used by previous works [9] where some images for each logo-class are utilized for training while the rest is used for evaluation. Despite our one-shot architecture, we evaluated the framework in the traditional setup in order to directly measure our model's performance against the traditional setting that has been adopted by every existing work. Although different classification and recognition problems have been extended to

query-based retrieval setups (e.g. image classification to image retrieval [61], handwriting recognition to query-based word spotting [62]), not many prior works using deep learning frameworks exist that addresses logo retrieval in an open-vocabulary scenario and is robust to challenges like low resolution, cluttered background etc. While existing methods are designed for more complicated tasks (e.g. detection and localization of all the logos in the training set), the boost in performance even in sparse training data plus the ability to generalize to unseen logos in the wild justifies our proposal to extend the problem-statement from a traditional to query-based setup.

### 5.1. Datasets

Our model makes use of binary masks as ground truths. FlickrLogos-32 [63] is a very popular logo dataset consisting of boundary box annotations as well as binary masks. Therefore, we train our model on this dataset. To evaluate the robustness of the one-shot architecture we explored another dataset TopLogos-10 [9]. The main reasons to select FlickrLogos for training and TopLogos for testing are as follows: (a) To the best of our knowledge, FlickrLogos is the only logo dataset which has binary segmentation mask as ground truth along with bounding box labels. Rest of the available datasets have bounding box labels only. (b) During testing, we preferred to use TopLogos because it has instances of logos of varying sizes and different cluttered background scenario which provides many challenges for logo detection. Brief discussions of these datasets are given below.

**Training set - FlickrLogos-32:** This dataset comprises 8240 images from 32 different logo classes, each class representing a particular brand. Each class has 70 images with ground truth annotations in the form of bounding boxes and binary masks. To make the dataset congenial for our approach, we ignore 6000 images with no logo class. Thus, we have total  $32 \times 70$  i.e. 2240 images available for our experiments.

**Evaluation set - TopLogos-10:** It contains 700 images of 10 different clothing brand logos with various degree of composition complexity in the logos. Basically, there are ten logo classes: Adidas, Chanel, Gucci, Helly Hansen, Lacoste, Michael Kars, Nike, Prade, Puma, Supreme. For every class, there are 70 images with fully manually labelled logo bounding boxes. But, unlike the FlickrLogos-32 it does not have binary mask annotations. This dataset contains natural images where logo instances in a variety of context, e.g. Hats, wallets, shoes, shown gels, lipsticks, spray, phone covers, eye glasses, jackets, peaked caps, T-shirts, and sign boards. In short, TopLogos-10 represents logo instances with varying sizes in natural real-world scenarios that provide real challenges for logo detection. The main reason behind using TopLogos-10 dataset is to check the generalization ability of our framework in a completely unseen scenario with more complexity, which is a pressing need of one-shot framework. Even if TopLogos has only 10 logo classes, every query-target pair (note that, this number is large) is unknown to the trained model during testing and it challenges the model to deal with real one-shot scenarios like varying logo sizes, different cluttered background etc.

### 5.2. Implementation details

We will discuss some salient details of our model: starting from data preparation to training and inference details.

**Data preparation:** (a) *One-shot setting:* As discussed earlier, we have used FlickrLogos-32 dataset for training our model in one-shot setting. It contains 2240 images with ground truth annotations. We denote this set as  $D^q$  and their corresponding masks as  $D^m$ . The images of these two sets are resized to a fixed dimension  $256 \times 256$ . Now, we obtain the conditional query logo objects by

cropping the main images with respect to their bounding boxes followed by resizing to a dimension of  $64 \times 64$ . This forms the conditional set of query objects  $D^q$ . Then, we generate a large number of triplets from the obtained image sets  $\{D^q, D^t, D^m\}$ . If  $n$  is the number of images per logo class then  $n \times (n - 1)$  triplet combination is possible per class. In total, for 32 logo classes we get  $32 \times (70 \times 69)$  i.e. 154,560 triplets. We use 90% of these triplets as our training set and 10% triplets as validation set. (b) *Traditional setting:* In this set-up, our main objective of query-based logo search remains unchanged. But, now we generate our training triplets only from a small set (40 images per class) of the available data in the FlickrLogos-32 dataset. These triplets will be used for training the model. And for testing we will use the rest of the images for each class.

**Training:** We train our model in an end-to-end fashion from scratch using these large number of generated triplets  $\{I_q^i, I_t^i, I_m^i\}_{i=1}^N$  where  $N$  is the total number of triplets. The conditioning branch takes  $I_q^i \in \mathbb{R}^{64 \times 64 \times 3}$  as an input and obtains a conditional representation of the logo object. This representation will be concatenated with the segmentation network at different scale. The segmentation branch takes  $I_t^i \in \mathbb{R}^{256 \times 256 \times 3}$  as input and tries to predict a binary segmentation mask  $I_m^i$ . We use Gaussian initialization with 0.01 standard deviation. We train our model using stochastic gradient descent (SGD) optimizer with initial learning rate 0.0004 and momentum and weight decay are set to 0.9 and 0.0005, respectively. The model is trained for 100K iterations with batch size of 32 for optimizing loss up to a satisfactory level. We have implemented the whole model in Tensorflow and run on a server with Nvidia Titan X GPU with 12 GB of memory.

**Testing:** In one-shot setting, the main advantage of our model is that once it's trained it can be used for any logo class. To test the scalability and robustness of our model we have used a different dataset. For traditional setting, we test our model on the same FlickrLogos-32 dataset but with testing images available for each class.

The final binary segmentation mask is obtained by using a threshold of 0.5 on the predicted mask. To obtain quantitative results on these datasets, we generate a minimum bounding box for each logo instances covering the mask from the predicted binary mask. To attain bounding boxes from the output map, first, we compute the topmost spatial position  $T(x_t, y_t)$ , bottom most binary pixel position  $B(x_b, y_b)$ , leftmost binary pixel position  $L(x_l, y_l)$  and topmost binary pixel position  $R(x_r, y_r)$ . Then, we evaluate bounding box value  $X' = x_l$ ,  $Y' = y_t$ ,  $H' = y_b - y_t$  and  $W' = x_r - x_l$ . At last, these bounding boxes are compared with the ground-truth bounding boxes to obtain the Intersection over Union (IoU). For the performance evaluation, we have used mean Average Precision (mAP) metric for all classes. For the mAP calculation, we choose the IoU threshold as 0.5. It means that a detection will be considered as positive if the IoU between the predicted and ground-truth exceeds 50%. As the proposed logo detection technique is based on segmentation of the logo, we have also included the mean pixel IoU or mPixIoU [64] as another metric for the evaluation purpose. However, unlike FlickrLogos-32 dataset, other dataset does not have any pixel level binary annotation. Thus, the mPixIoU metric is shown only for FlickrLogos-32 dataset. The code is available here<sup>1</sup>

### 5.3. Baselines

To exploit the robustness of our approach, we compared the following baselines with our proposed approach.

**Fine-tuning:** As suggested by Caelles et al. [21], we fine-tune a pre-trained FCN network (only the fully connected layers) with full

<sup>1</sup> <https://github.com/AyanKumarBhunia/Deep-One-Shot-Logo-Retrieval>.

**Table 1**

Comparison of Logo detection performance on FlickrLogos-32 dataset following traditional setting. “#” denotes query-based frameworks.

Method	mPixIoU	mIoU	mAP
Bag of Words (BoW) [41]	—	—	54.5
Deep Logo [13]	—	—	74.4
BD-FRCN-M [8]	—	—	73.5
Reallmg [9]	—	—	81.1
Faster-RCNN [4]	—	—	70.2
SSD [65]	—	—	67.5
YOLO [66]	—	—	68.7
U-Net [25]	20.3	30.1	40.5
#SiameseFCN [20]	70.8	77.6	79.4
#CoFCN [23]	71.4	81.9	84.1
#SG-One [24]	76.5	84.5	86.9
#Ours	<b>78.2</b>	<b>86.7</b>	<b>89.2</b>

supervision on the available paired data of a particular logo class and test it on the target images.

*U-Net*: Here, we focus on the naive way for generating segmentation maps by considering each logo class as a separate semantic class. We use U-Net [25] encoder-decoder network that takes the target image as input and outputs segmentation map without any conditional reference. Note that this set-up cannot be extended to new classes.

*SiameseFCN*: Siamese networks are extensively used for one-shot classification tasks [20]. In this method, we use two pre-trained FCNs to extract dense features from query images and target images. L1 similarity metric is used to learn the coherence between the feature for every pixel in the query image and the target image in order to produce a pixel-level binary mask.

*CoFCN*: Here, we explore a conditional segmentation network based on FCN [23]. At first, we feed query images in the conditional branch to generate a set of parameters  $\theta$ . We use  $\theta$  in the parameterized part of the learned segmentation model, which takes target images as input and produces a segmentation map.

*SG-One*: This one-shot segmentation framework [24] uses a new masked average pooling operation to extract the latent representation of the query image. However, since our framework uses cropped query image (i.e. support set), we use simple global average pooling instead, and use cosine-similarity followed by a tanh layer for similarity matching with rest setup similar to [24]. A quite similar one-shot framework has also been addressed for visual tracking problem [24].

#### 5.4. Evaluations

*Traditional setting*: We have compared our results with the baselines and some state-of-the-art logo detection methods [8,9,13,41] as well as evaluated the performance using popular bounding box detectors like Faster-RCNN<sup>2</sup> [4], YOLO<sup>3</sup> [66] and SSD<sup>4</sup> [65] while opting for VGG as the baseline architecture. Due to limited data, we initialize this framework with weights trained on Pascal-VOC dataset (see Table 1). Among popular traditional logo detectors, Reallmg [9] gives highest mAP score of 81.1% and uses a deep model to synthetically generate training data in order to improve its performance. Though state-of-the-art object detectors like Faster-RCNN, YOLO, SSD achieve impressive performance on benchmark object detection datasets, their performance is limited in FlickrLogos-32 dataset because of limited training data. For the same reason, the performance of U-Net is poor. SiameseFCN, which is a query-based framework, performed moderately

**Table 2**

Logo detection performance on Evaluation set following one-shot setting.

Training set	Testing set	Method	mPixIoU	mIoU	mAP
FlickrLogos (20 classes)	FlickrLogos (12 classes)	Fine-tuning	15.8	24.1	27.4
		SiameseFCN	36.9	46.5	51.1
		CoFCN	44.6	51.3	59.7
		SG-One	52.3	59.7	64.4
		<b>Ours</b>	<b>59.1</b>	<b>66.3</b>	<b>66.8</b>
FlickrLogos (20 classes)	TopLogos (10 classes)	Fine-tuning	—	—	14.7
		SiameseFCN	—	—	24.4
		CoFCN	—	—	29.2
		SG-One	—	—	36.9
		<b>Ours</b>	—	—	<b>40.1</b>

on FlickrLogo-32 dataset. CoFCN performed well with a high mAP of 84.1% which is 3.0% and 9.4% greater than Reallmg [9] and Deep Logo [13] respectively. The proposed multi-scale query based detection technique achieved a better mAP than the previous state-of-the-art method [9] with 8.1% increase in absolute mAP; our rival query-based method, SG-One [24] trails by 2.3% mAP value.

*One-shot setting*: For the one-shot setup, the results shown in Table 2 indicate the capability of different methods to generalize to new classes. These results are comparatively lower than the results observed in the traditional setting. We realize that this is because we are trying to evaluate the performance of the model on a more open space setting by allowing the system to detect logo in a target scene from a single reference logo sample; contrarily, in the traditional setup, the total number of logo classes is fixed and 40 images from each logo class were used to train the model. The fine-tuned baseline produces relatively low mAP since it quickly overfits to the fine-tuning data in the support set. For FlickrLogo-32, our proposed method outperformed the general fine-tuning based approach by 39.4% mAP value. On TopLogos-10, which is a clothing-logo dataset with dense background clutter, observed results are relatively low. We demonstrate some qualitative results of our framework in Fig. 3. We have also evaluated the performance by varying IoU threshold from 0.5 to 0.8 and the performance of our framework drops to a limited extent compared to other competitive methods (see Fig. 4), and thus it illustrates the superiority of our design choices.

These improvements over other existing one-shot frameworks [20,23,24] are mainly due to following design choices in our framework. (1) Multi-scale conditioning (2) Parameterized similarity matching operation through  $1 \times 1$  convolutional layer at every scale (3) Skip connection between encoder and decoder part of the network. Overall, our framework is undoubtedly simple, easy to implement and understand.

#### 5.5. Ablation study

In this section, we have given a comprehensive analysis of each sub-variant of our framework along with quantitative analysis for different alternative network design choices. Here, we have also shown results for  $k$ -shot cases, using  $k$  different samples of the query image as our conditional input serially, and final binary mask is obtained by logical OR operation between the  $k$  predicted binary maps. For the traditional set-up we have used FlickrLogos-32 dataset as before. For  $k$ -shot setting, we have trained our model using images of 20 classes of FlickrLogos-32 dataset and the remaining part is used for testing (Table 3).


- Variant 1: Here, we use a general U-Net architecture without any conditioning branch(not a query-based framework) and notice limited performance.
- Variant 2: In this variant, we use conditional query representation, which is concatenated with the latent space of segmentation module, in order to fetch of the location

<sup>2</sup> <https://goo.gl/zQsSka> (Faster-RCNN).


<sup>3</sup> <https://goo.gl/vUFEuq> (YOLO).

<sup>4</sup> <https://goo.gl/vxR9c7> (SSD).






**Fig. 3.** Some qualitative results of our method: In the first two rows we have shown results on FlickrLogo-32 using traditional setting and for the next two rows we have shown one-shot results on TopLogos-10 dataset. Here, images are in the following order (left to right): query logo, target image, predicted mask, and final detection result (shown by boxes on the images), respectively.



**Fig. 4.** Performance with varying IoU threshold values.

**Table 3**  
mAP (%) of different variants of our configuration on FlickrLogos-32 dataset using traditional setting.

Method	<i>k</i> -shot setting			Traditional Setting
	<i>k</i> =1	<i>k</i> =3	<i>k</i> =5	
Variant 1	—	—	—	36.5
Variant 2	59.1	64.5	65.6	82.7
<b>Proposed</b>	<b>66.8</b>	<b>71.2</b>	<b>71.6</b>	<b>89.2</b>



**Fig. 5.** Some failure cases of our frameworks. Yellow box is the Ground-Truth. Green box is True Positive. Red box is False Positive. First row illustrates a false-positive case where our model wrongly predicts logo instances (marked in red). This happens due to certain extent of similarity between the query logo and irrelevant patch in the target image. In the second row, it fails to detect a logo instances due to different orientation. Third row shows a failure case due to tiny size of the logo and irregular illumination. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.).

of query logo in a target image. However, it fails to detect small logos in images, specifically when the size of the query logo image and the actual logo present in the target differs considerably.

- **Proposed Method:** We overcome the above problems by employing a *multi-scale conditioning* operation. This makes our model capable of detecting small transformed logos, which is the one of the major challenges for logo detection.

**Alternative designs:** To localize the query logo within the target scene image, instead of cosine-similarity measure, we focus on simple concatenation operation followed by  $1 \times 1$  convolution to learn the correlation. Using a parameterized layer for similarity matching at multiple-scales offers a better performance in the context of logo detection (see Section 4.3). Keeping rest of the architecture same, we calculate cosine similarity at each spatial position of feature-map from encoder (segmentation) network with conditional query latent vector (converted to the same depth as  $f^{i,s}$  using  $1 \times 1$  convolution), followed by a *tanh* layer and finally combine to the corresponding decoder stage. This alternative network design achieves mAP values of 64.1% and 87.8%, trailing ours by 2.7% and 1.4% for one-shot and traditional setting on the FlickrLogos-32 dataset (same setting as Table 3). Also, we tried using pooled (Global Average Pooling [67]) feature representation of feature maps from multiple layers of the conditional network, with an intuition to get a multi-scale representation of the query logo and evaluate similarity (*ours*) with feature maps of same stage from segmentation network; the results, however, drop by 5.6%. This signifies that the final latent representation from deep layers of the conditional network contains higher semantic information of the logo that can be used as a conditional input to segmentation network to fetch the logo at multiple-scales of the target image.

**Error analysis:** Fig. 5 shows some of the failure cases of our framework. Here, we have used simple pixel-wise cross-entropy loss to train the network; however, rotation invariant network design could be explored on the top of our proposed multi-scale conditioning framework, so that rotation related error (see second row of 5) could be avoided. Adversarial feature deformation



[68] could be used to generalize the model from limited data. In addition, it is to be noted that in spite of multiscale-conditioning it fails in few cases for tiny size of the logo. To alleviate such errors, a feature magnification module like [69] or attention mechanism [70] could be helpful. In addition, the imbalance between foreground-background could be better handled using newly proposed focal-loss [71] which adds a multiplying factor to the normal cross-entropy loss in order to emphasize on hard, misclassified examples.

## 6. Conclusion

In this paper, we have re-designed the traditional logo detection problem setting by proposing a query based logo search system that uses a one-shot architecture. The driving idea of our architecture is the use of *multi-scale conditioning* with a skip-connection based architecture that predicts a logo segmentation mask. The proposed framework is simple and easy to implement. It is capable of detecting new logo classes without additional training data.

We demonstrate the effectiveness of our system by doing experiment on unseen logos in the wild. From the experiments on publicly available logo detection datasets, we noted that our proposed system outperformed the benchmark results [8,9,13,41] without even extending existing logo datasets.

Though our system is scale invariant, it may fail for logos which are tiny in size. A feature magnification module [69] could be useful to improve the performance in such cases. Also, the imbalance between fore-ground and background information could be handled properly to boost the performance further. In future we plan to work on these to improve the accuracy.

## References

- [1] Y. Gao, F. Wang, H. Luan, T.S. Chua, Brand data gathering from live social media streams, in: Proceedings of the ICMR, ACM, 2014, p. 169.
- [2] A.P. Psyllos, C.-N.E. Anagnostopoulos, E. Kayafas, Vehicle logo recognition using a sift-based enhanced matching scheme, IEEE Trans. Intell. Transp. Syst. 11 (2) (2010) 322–328.
- [3] N. Hagbi, O. Bergig, J. El-Sana, M. Billinghurst, Shape recognition and pose estimation for mobile augmented reality, IEEE Trans. Visual. Comput. Graph. 17 (10) (2011) 1369–1379.
- [4] S. Ren, K. He, R. Girshick, J. Sun, Faster R-CNN: Towards real-time object detection with region proposal networks, in: Proceedings of the Advances in Neural Information Processing Systems, 2015, pp. 91–99.
- [5] S. Bianco, M. Buzzelli, D. Mazzini, R. Schettini, Deep learning for logo recognition, Neurocomputing 245 (2017) 23–30.
- [6] R. Boia, C. Florea, L. Florea, R. Dogaru, Logo localization and recognition in natural images using homographic class graphs, Mach. Vis. Appl. 27 (2) (2016) 287–301.
- [7] C. Eggert, A. Winschel, R. Lienhart, On the benefit of synthetic data for company logo detection, in: Proceedings of the ACM Multimedia, ACM, 2015, pp. 1283–1286.
- [8] G. Oliveira, X. Frazão, A. Pimentel, B. Ribeiro, Automatic graphic logo detection via fast region-based convolutional networks, in: Proceedings of the IJCNN, IEEE, 2016, pp. 985–991.
- [9] H. Su, X. Zhu, S. Gong, Deep learning logo detection with data expansion by synthesising context, in: Proceedings of the WACV, 2017, pp. 530–539.
- [10] A. Alaei, P.P. Roy, U. Pal, Logo and seal based administrative document image retrieval: a survey, Comput. Sci. Rev. 22 (2016) 47–63.
- [11] G. Csurka, C. Dance, L. Fan, J. Willamowski, C. Bray, Visual categorization with bags of keypoints, in: Proceedings of the ECCV, Vol. 1, Prague, 2004, pp. 1–2.
- [12] S. Bianco, M. Buzzelli, D. Mazzini, R. Schettini, Logo recognition using CNN features, in: Proceedings of the ICIA, Springer, 2015, pp. 438–448.
- [13] F.N. Iandola, A. Shen, P. Gao, K. Keutzer, Deeplogo: hitting logo recognition with the deep neural network hammer, arXiv:1510.02131, (2015).
- [14] R. Girshick, J. Donahue, T. Darrell, J. Malik, Rich feature hierarchies for accurate object detection and semantic segmentation, in: Proceedings of the CVPR, 2014, pp. 580–587.
- [15] R. Girshick, Fast R-CNN, in: Proceedings of the ICCV, 2015, pp. 1440–1448.
- [16] Y. Liu, L. Jin, S. Zhang, C. Luo, S. Zhang, Curved scene text detection via transverse and longitudinal sequence connection, Pattern Recogn. 90 (2019) 337–345.
- [17] C. Luo, L. Jin, Z. Sun, Moran: a multi-object rectified attention network for scene text recognition, Pattern Recogn. 90 (2019) 109–118.
- [18] A.K. Bhunia, A. Konwer, A.K. Bhunia, A. Bhowmick, P.P. Roy, U. Pal, Script identification in natural scene image and video frames using an attention based convolutional-LSTM network, Pattern Recogn. 85 (2019) 172–184.
- [19] L. Fei-Fei, R. Fergus, P. Perona, One-shot learning of object categories, IEEE Trans. Pattern Anal. Mach. Intell. 28 (4) (2006) 594–611.
- [20] G. Koch, R. Zemel, R. Salakhutdinov, Siamese neural networks for one-shot image recognition, in: Proceedings of the ICML, Vol. 2, 2015.
- [21] S. Caelles, K.-K. Maninis, J. Pont-Tuset, L. Leal-Taixé, D. Cremers, L. Van Gool, One-shot video object segmentation, in: Proceedings of the CVPR, IEEE, 2017.
- [22] N. Dong, E.P. Xing, Few-shot semantic segmentation with prototype learning, in: Proceedings of the BMVC, Vol. 3, 2018, p. 4.
- [23] A. Shaban, S. Bansal, Z. Liu, I. Essa, B. Boots, One-shot learning for semantic segmentation, Proceedings of the BMVC, 2017 arXiv:1709.03410.
- [24] X. Zhang, Y. Wei, Y. Yang, T. Huang, Sg-one: similarity guidance network for one-shot semantic segmentation, 2018 arXiv:1810.09091.
- [25] O. Ronneberger, P. Fischer, T. Brox, U-net: Convolutional networks for biomedical image segmentation, in: Proceedings of the MICCAI, 2015, pp. 234–241.
- [26] Y. Lamdan, J.T. Schwartz, H.J. Wolfson, Object recognition by affine invariant matching, in: Proceedings of the CVPR, IEEE, 1988, pp. 335–344.
- [27] C. Schmid, R. Mohr, Local Grayvalue invariants for image retrieval, IEEE Trans. Pattern Anal. Mach. Intell. 19 (5) (1997) 530–535.
- [28] T. Gevers, A.W. Smeulders, Color-based object recognition, Pattern Recogn. 32 (3) (1999) 453–464.
- [29] D.G. Lowe, Distinctive image features from scale-invariant keypoints, Int. J. Comput. Vis. 60 (2) (2004) 91–110.
- [30] Y.-S. Kim, W.-Y. Kim, Content-based trademark retrieval system using a visually salient feature, Image Vis. Comput. 16 (12–13) (1998) 931–939.
- [31] P.-Y. Yin, C.-C. Yeh, Content-based retrieval from trademark databases, Pattern Recogn. Lett. 23 (1–3) (2002) 113–126.
- [32] J.P. Eakins, K.J. Riley, J.D. Edwards, Shape feature matching for trademark image retrieval, in: Proceedings of the ICIVR, Springer, 2003, pp. 28–38.
- [33] R.H. Van Leuken, M.F. Demirci, V.J. Hodge, J. Austin, R.C. Veltkamp, Layout indexing of trademark images, in: Proceedings of the ICIVR, ACM, 2007, pp. 525–532.
- [34] B. Kovar, A. Hanjalic, Logo appearance statistics in a sport video: Video indexing for sponsorship revenue control, in: Proceedings of the SPIE, 2002.
- [35] R.J. Den Hollander, A. Hanjalic, Logo recognition in video stills by string matching, in: Proceedings of the ICIP, Vol. 3, IEEE, 2003, pp. III–517.
- [36] D. Hall, F. Pélissier, O. Riff, J.L. Crowley, Brand identification using gaussian derivative histograms, Mach. Vis. Appl. 16 (1) (2004) 41–46.
- [37] A.D. Bagdanov, L. Ballan, M. Bertini, A. Del Bimbo, Trademark matching and retrieval in sports video databases, in: Proceedings of the Multimedia Information Retrieval, ACM, 2007, pp. 79–86.
- [38] J. Kleban, X. Xie, W.-Y. Ma, Spatial pyramid mining for logo detection in natural scenes, in: Proceedings of the ICMR, IEEE, 2008, pp. 1077–1080.
- [39] A. Joly, O. Buisson, Logo retrieval with a contrario visual query expansion, in: Proceedings of the ACM Multimedia, ACM, 2009, pp. 581–584.
- [40] J. Meng, J. Yuan, Y. Jiang, N. Narasimhan, V. Vasudevan, Y. Wu, Interactive visual object search through mutual information maximization, in: Proceedings of the ACM Multimedia, ACM, 2010, pp. 1147–1150.
- [41] S. Romberg, R. Lienhart, Bundle min-hashing for logo recognition, in: Proceedings of the ICMR, ACM, 2013, pp. 113–120.
- [42] S. Romberg, L.G. Pueyo, R. Lienhart, R. Van Zwol, Scalable logo recognition in real-world images, in: Proceedings of the ICMR, ACM, 2011, p. 25.
- [43] J. Revaud, M. Douze, C. Schmid, Correlation-based burstiness for logo retrieval, in: Proceedings of the ACM Multimedia, ACM, 2012, pp. 965–968.
- [44] R. Boia, C. Florea, L. Florea, Elliptical asifit agglomeration in class prototype for logo detection, in: Proceedings of the BMVC, 2015, pp. 115–1.
- [45] M. Rusiñol, J. Lladós, Efficient logo retrieval through hashing shape context descriptors, in: Proceedings of the DAS, ACM, 2010, pp. 215–222.
- [46] J. Fu, J. Wang, H. Lu, Effective logo retrieval with adaptive local feature selection, in: Proceedings of the ACM Multimedia, ACM, 2010, pp. 971–974.
- [47] J. Wang, Q. Liu, J. Liu, H. Lu, Logo retrieval with latent semantic analysis, in: Proceedings of the VIP, Citeseer, 2006.
- [48] M. Soysal, A.A. Alatan, Joint utilization of appearance, geometry and chance for scene logo retrieval, Comput. J. 54 (7) (2011) 1221–1231.
- [49] H. Su, X. Zhu, S. Gong, Open logo detection challenge, Proceedings of the BMVC, 2018 arXiv:1807.01964.
- [50] H. Su, S. Gong, X. Zhu, Scalable deep learning logo detection, arXiv:1803.11417, (2018b).
- [51] O. Vinyals, C. Blundell, T. Lillicrap, D. Wierstra, et al., Matching networks for one shot learning, in: Proceedings of the Advances in neural information processing systems, 2016, pp. 3630–3638.
- [52] L. Bertinetto, J.F. Henriques, J. Valmadre, P. Torr, A. Vedaldi, Learning feed-forward one-shot learners, in: Proceedings of the Advances in Neural Information Processing Systems, 2016, pp. 523–531.
- [53] B. Hariharan, R. Girshick, Low-shot visual recognition by shrinking and hallucinating features, in: Proceedings of the ICCV, 2017.
- [54] H. Noh, P. Hongsuck Seo, B. Han, Image question answering using convolutional neural network with dynamic parameter prediction, in: Proceedings of the CVPR, 2016.
- [55] N. Dong, E.P. Xing, Few-shot semantic segmentation with prototype learning, in: Proceedings of the BMVC, Vol. 3, 2018, p. 4.
- [56] S. Bell, C. Lawrence Zitnick, K. Bala, R. Girshick, Inside-outside net: Detecting objects in context with skip pooling and recurrent neural networks, in: Proceedings of the CVPR, 2016, pp. 2874–2883.
- [57] B. Hariharan, P. Arbeláez, R. Girshick, J. Malik, Hypercolumns for object segmentation and fine-grained localization, in: Proceedings of the CVPR, 2015, pp. 447–456.

- [58] T. Kong, A. Yao, Y. Chen, F. Sun, Hypernet: Towards accurate region proposal generation and joint object detection, in: Proceedings of the CVPR, 2016, pp. 845–853.
- [59] H. Wang, Q. Wang, M. Gao, P. Li, W. Zuo, Multi-scale location-aware kernel representation for object detection, in: Proceedings of the CVPR, 2018, pp. 1248–1257.
- [60] K. He, G. Gkioxari, P. Dollár, R. Girshick, Mask R-CNN, in: Proceedings of the ICCV, IEEE, 2017, pp. 2980–2988.
- [61] A. Gordo, J. Almazán, J. Revaud, D. Larlus, Deep image retrieval: Learning global representations for image search, in: Proceedings of the ECCV, Springer, 2016, pp. 241–257.
- [62] R. Manmatha, C. Han, E.M. Riseman, Word spotting: a new approach to indexing handwriting, in: Proceedings of the CVPR, IEEE, 1996, pp. 631–637.
- [63] S. Romberg, L.G. Pueyo, R. Lienhart, R. van Zwol, Scalable logo recognition in real-world images, in: Proceedings of the ICMR, 2011, pp. 25:1–25:8.
- [64] C. Liang-Chieh, G. Papandreou, I. Kokkinos, K. Murphy, A. Yuille, Semantic image segmentation with deep convolutional nets and fully connected crfs, in: Proceedings of the ICLR, 2015.
- [65] W. Liu, D. Anguelov, D. Erhan, C. Szegedy, S. Reed, C.-Y. Fu, A.C. Berg, SSD: Single shot multibox detector, in: Proceedings of the ECCV, Springer, 2016, pp. 21–37.
- [66] J. Redmon, S. Divvala, R. Girshick, A. Farhadi, You only look once: Unified, real-time object detection, in: Proceedings of the CVPR, 2016, pp. 779–788.
- [67] M. Lin, Q. Chen, S. Yan, Network in network, Proceedings of the ICLR, 2014 arXiv:1312.4400.
- [68] A.K. Bhunia, A. Das, A.K. Bhunia, P. Sai Raj Kishore, P. Pratim Roy, Handwriting recognition in low-resource scripts using adversarial learning, in: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, 2019, pp. 4767–4776.
- [69] J. Li, X. Liang, Y. Wei, T. Xu, J. Feng, S. Yan, Perceptual generative adversarial networks for small object detection, in: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, 2017, pp. 1222–1230.
- [70] A.K. Bhunia, S.R.K. Perla, P. Mukherjee, A. Das, P.P. Roy, Texture synthesis guided deep hashing for texture image retrieval, in: Proceedings of the IEEE Winter Conference on Applications of Computer Vision (WACV), IEEE, 2019, pp. 609–618.
- [71] T.-Y. Lin, P. Goyal, R. Girshick, K. He, P. Dollár, Focal loss for dense object detection, in: Proceedings of the IEEE International Conference on Computer Vision, 2017, pp. 2980–2988.




**Ayan Kumar Bhunia** completed B.Tech in 2017 from Institute of Engineering & Management, Kolkata, India. He is now pursuing PhD at Centre for Vision, Speech and Signal Processing (CVSSP) of the University of Surrey, United Kingdom. Prior to that, he worked as a research assistant at Institute for Media Innovation (IMI) Lab of Nanyang Technological University, Singapore.



**Ankan Kumar Bhunia** is pursuing B.E. from Jadavpur University, Kolkata, India.



**Shuvojit Ghose** is pursuing B.Tech from Institute of Engineering & Management, Kolkata, India.



**Abhirup Das** is pursuing B.Tech from Institute of Engineering & Management, Kolkata, India.



**Partha Pratim Roy** received his Ph.D. degree in Computer Science in 2010 from Universitat Autònoma de Barcelona, (Spain). He worked as postdoctoral research fellow in the Computer Science Laboratory (LI, RFAI group), France and in Synchronmedia Lab, Canada. Presently, Dr. Roy is working as assistant professor at Indian Institute of Technology (IIT), Roorkee. His main research area is Pattern Recognition.



**Umapada Pal** received his Ph.D. from Indian Statistical Institute. He did his Post-Doctoral at INRIA (Institut National de Recherche en Informatique et en Automatique), France. From January 1997, he is a Faculty member of the Computer Vision and Pattern Recognition Unit (CVPRU), Indian Statistical Institute, Kolkata and at present he is a Professor and Head of CVPRU. Because of his significant impact in the Document Analysis research domain of Indian language, TC-10 and TC-11 committees of IAPR (International Association for Pattern Recognition) presented 'ICDAR Outstanding Young Researcher Award' to Dr. Pal in 2003. He is the Editorial board member several journals like PR, IJDAR, PRL, ACM Transactions on Asian Language Information Processing etc. He is a fellow of IAPR and senior member of IEEE.