

Sketch-Guided Object Localization in Natural Images

Aditay Tripathi¹, Rajath R Dani¹, Anand Mishra² and Anirban Chakraborty¹

¹Indian Institute of Science, Bengaluru ²Indian Institute of Technology, Jodhpur
<http://visual-computing.in/sketch-guided-object-localization/>

Abstract. We introduce a novel problem of localizing all the instances of an object (seen or unseen during training) in a natural image via sketch query. We refer to this problem as sketch-guided object localization. This problem is distinctively different from the traditional sketch-based image retrieval task where the gallery set often contains images with only one object. The sketch-guided object localization proves to be more challenging when we consider the following: (i) the sketches used as queries are abstract representations with little information on the shape and salient attributes of the object, (ii) the sketches have significant variability as they are hand-drawn by a diverse set of untrained human subjects, and (iii) there exists a domain gap between sketch queries and target natural images as these are sampled from very different data distributions. To address the problem of sketch-guided object localization, we propose a novel *cross-modal attention* scheme that guides the region proposal network (RPN) to generate object proposals relevant to the sketch query. These object proposals are later scored against the query to obtain final localization. Our method is effective with as little as a single sketch query. Moreover, it also generalizes well to object categories not seen during training and is effective in localizing multiple object instances present in the image. Furthermore, we extend our framework to a multi-query setting using novel feature fusion and attention fusion strategies introduced in this paper. The localization performance is evaluated on publicly available object detection benchmarks, viz. MS-COCO and PASCAL-VOC, with sketch queries obtained from ‘Quick, Draw!’. The proposed method significantly outperforms related baselines on both single-query and multi-query localization tasks.

Keywords: Sketch, Cross-modal retrieval, Object localization, One-shot learning, Attention

1 Introduction

Localizing objects in a scene via an image query has been a long-sought pursuit in computer vision literature. The seminal paper Video Google [29] addresses object localization in videos using a text retrieval approach. More recent works




Fig. 1: **Sketch-guided object localization.** Can we localize a cake, a pizza and birds in these images by just drawing their sketches? In this paper, we introduce the problem of sketch-guided object localization, and propose a novel *cross-modal attention* scheme to address this problem. **[Best viewed in color].**

on this problem focus on localizing seen as well as unseen object categories with as little as one image query during test time [12,17]. There are several applications associated with object localization in an image via image query - examples include object tracking [17] and content-based image retrieval [27,33]. However, the image of an object may not always be preferred as a query due to many practical reasons. These include (i) copyright issues, (ii) privacy concerns, or (iii) significant overhead to collect and annotate images of a rapidly expanding object category set, e.g., in industrial applications where images of each part of the equipment may not be available during the training. Further, it is also not practical to assume that the object names are always available for localization. In such situations, where stronger visual or semantic cues are unavailable, it would be worth exploring if a hand-drawn sketch of an object can be utilized for localizing the object in a natural image. In other words, given hand-drawn sketches of objects (for example, cake, pizza, and bird as shown in the top row of Fig. 1) can we design a framework that learns to localize all instances of these objects present in natural images? In this work, we explore this problem for the first time in the literature and propose a solution to it.

Despite the problem of sketch-guided object localization has a relation to the well-studied sketch-based image retrieval (SBIR) problem [5,14,20,25,28,30,38,40], these are distinctively different in terms of objectives. It is important to note that SBIR aims to retrieve images from a gallery of localized objects for a given sketch query. Contrary to this, our objective is to precisely localize all instances of an object of interest in a natural scene in presence of many other distracting objects. Furthermore, sketch-guided object localization poses several additional challenges compared to image query-guided object localization, notably - (i) the sketches used as queries are abstract drawings with little information on shape and salient attributes of the object, (ii) the sketches have significant variability as they are drawn by a diverse set of untrained human subjects, and (iii) a large domain gap exists between the sketch queries and the target gallery images.

Plausible baseline approaches vs. proposed solution: Despite its practical utility, sketch-guided object localization has never been explored in the literature. In this work, we consider a state-of-the-art image query-guided localization method [12] and a modified Faster R-CNN [24] as probable baseline approaches. We empirically show that these methods are insufficient due to their sensitivity to the domain gap present between sketches and natural images, and their ineffective attention mechanisms in this context. To mitigate these issues, we develop a *cross-modal attention* scheme to guide proposal generation, which is a crucial step in the localization framework. The proposed *cross-modal attention* scheme generates a spatial compatibility matrix by comparing the global sketch representation with the local image representations obtained from each location of the image feature map. This spatial compatibility matrix creates an attention feature map that is fused with the original feature map of the image. The result of this fusion is subsequently fed to a region proposal network (RPN) to generate proposals relevant to the sketch query. Finally, the proposals are pooled and compared against the query to precisely localize the object of interest.

Our proposed model, by virtue of the aforementioned *cross-modal attention*, embeds query information in the image feature representation before generating region proposals. Therefore, it is inherently able to generate relevant object proposals even for object categories unseen during the training time. This enables the proposed method to achieve superior performance to baseline methods for unseen categories while achieving reasonably high performance for seen categories as well. We also explore the possibility of using multiple sketch queries towards the localization task. To generalize our proposed model to the multi-query setting, we propose novel feature and attention pooling mechanisms.

Contributions: In summary, we make the following contributions:

- (i) We introduce an important but unexplored problem – *sketch-guided object localization in natural images*. This novel problem is well motivated from scenarios where a query image corresponding to the object of interest or the object category names are not available, but a sketch can be hand-drawn to provide a minimal visual description of the object.
- (ii) We propose a novel *cross-modal attention* scheme to guide the region proposal network to generate object proposals relevant to the sketch query. Our method is effective with as little as one sketch used as a query. Moreover, it generalizes well to unseen object categories with $\approx 3\%$ improvement over a one-shot detection method used as baseline [12] and is also effective in localizing multiple object instances present in the image.
- (iii) To support multiple sketch-query based object localization, we propose feature and attention pooling in our architecture and demonstrate promising performance.
- (iv) We have performed extensive experimentation and ablation studies on two public benchmarks. We firmly believe that our work will open-up novel future research directions under sketch-guided computer vision tasks.

2 Related work

Sketch-based Image Retrieval: Given a sketch query, sketch-based image retrieval (SBIR) aims to retrieve from a gallery of images containing a single object. Methods in SBIR can be grouped into two broad categories- (a) *classical methods* which use hand-crafted features, for example, SIFT or gradient-field histogram of gradients along with a bag-of-words representation of sketches as shown in [5,14], (b) *deep methods* which utilize cross-modal deep learning techniques by incorporating ranking loss such as the contrastive loss [25] or the triplet loss [38] to learn a ranking function between candidate images and sketch queries which are then subsequently used to score the candidate images, and the top-scoring image is retrieved. In [30], researchers have utilized an attention model to solve fine-grained SBIR, and have introduced HOLEF loss to bridge the domain gap between the sketches and images. Since the time complexity of such methods for large-scale SBIR problem is significantly high, researchers have proposed hashing models, such as [20,28,35,40], which significantly reduces the retrieval time. Contrary to these works, our goal is to use a sketch query to *localize* all instances of an object in an image that usually contains many objects of different categories.

Object Detection: Object detection involves localizing and classifying an object in a given image. State-of-the-art object localization methods can be primarily grouped into two categories: *proposal-free* [16,18,21,22,23,26] and *proposal-based* [1,7,8,10,11,24]. *Proposal-free* methods are single-stage detectors, and are faster during inference. However, they often fall short in performance as compared to *proposal-based* methods. *Proposal-based* methods first generate object proposals and then refine them by classifying each of them into one of the object categories. In [8], one such method, selective search [32] was utilized to generate proposals in the first stage, and then use these generated proposals to classify the object. However, the two-stages in this model were trained independently. Faster R-CNN [24] introduced a region proposal network (RPN) that made the detection pipeline end-to-end trainable. Note that, both these types of object detectors are query-free. More recently, Hsieh et al. [12] have introduced one-shot object detection via an image query, where the goal is to localize all the instances of an unseen object in the target image via an image query of the same object. Unlike their work, where query and target images are from the same distribution, our queries, i.e., hand-drawn sketches, are from a significantly different domain as compared to that of the target images.

Attention in Deep Neural Networks: The attention model in deep neural networks helps the salient features of the input image to become more critical when required. In [3,31], the authors proposed attention networks to generate attention scores on each of the object proposals, and an attention-based dropout layer for weakly supervised object localization. In [2] a dynamic attention-action strategy using deep reinforcement learning was proposed to localize objects in an image. Hsieh et al. [12] adapted the self-attention mechanism proposed in [34]

and applied it to image query-guided object detection. In this method, each pixel in an image was represented as a weighted combination of the pixels in the query image. The weights depend on the similarity between each image pixel and query pixel pairs for all pixels in the query. Conversely, our proposed *cross-modal attention* computes the spatial compatibility between global query representation and localized representations of image regions. This enables it to mitigate the domain misalignment problem prevalent in sketch-guided object localization. Yan et al. [37] used class-specific attentive vectors inferred from images of objects in a meta-set for applying channel-wise soft attention to feature maps of the proposals. However, the channel-wise soft attention may not be trivially utilized in our case due to the domain gap between query sketch and target natural images.

Sketch Representation: Besides the traditionally used convolutional neural networks [39], there have been some recent advancements towards robust representation learning for sketches, e.g., SketchRNN [9], transformer-based architecture [36]. Although we have chosen to utilize CNN (ResNet)-based feature encoders in this work, our proposed localization framework can seamlessly integrate more advanced sketch representation learning methods in the feature extraction modules.

3 Proposed Methodology

In this section, we first formally introduce the novel problem of sketch-guided object localization in natural images. Then, we present a solution framework built around our novel *cross-modal attention* scheme and discuss its utility for one-shot object detection using a sketch query. Further, we extend our proposed scheme to a multi-query setting, i.e., *multiple sketch query-based object localization* task by introducing principled fusion strategies.

3.1 Problem Formulation

Let $I = \{I_{train}, I_{test}\}$ be a set of all natural images, each containing variable number of object instances and categories, in a dataset \mathcal{D}_I . Here I_{train} and I_{test} are sets of train and test images respectively. Like any other machine learning task these two sets are mutually exclusive and only I_{train} is available during training time. Further, let $S = \{S_{train}, S_{test}\}$ be a set of all sketches, each containing one object, and $C = \{C_{train}, C_{test}\}$ be a set of all object categories. During training time, each training sample contains an image $i \in I_{train}$, a query sketch $s_c \in S_{train}$, where $c \in C_{train}$, and all the bounding boxes corresponding to object category c in the image i . During test time, given an image $i' \in I_{test}$ and sketch query $s_c \in S_{test}$, where $c \in C_{test}$, the problem is to localize all the instances of the object category c in the image i' . Note that we show experimental results in cases where $C_{train} = C_{test}$ i.e., categories in C_{test} are seen during training time (common train-test categories), as well as $C_{train} \cap C_{test} = \phi$, i.e., categories in C_{test} are not seen during training (disjoint train-test categories).

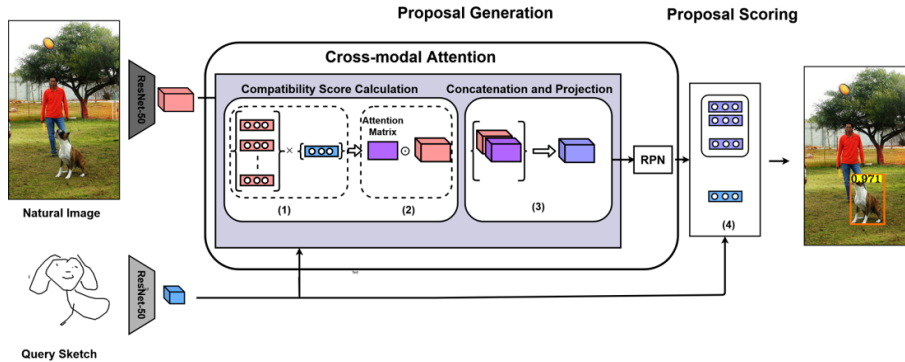


Fig. 2: Given an image and a query sketch, our end-to-end trainable sketch-guided localization framework works in the following two stages: (i) **query-guided proposal generation:** in this stage, feature vectors corresponding to different regions in the image feature map (shown using pink) are scored with the global sketch representation (shown using blue) to identify the compatibility (Block-1). Then, these compatibility scores (shown using violet) are multiplied with image feature maps (shown using pink) to get attention feature (Block-2). Further, these attention feature maps are concatenated with the original feature maps and projected to lower-dimensional space which is then passed through region proposal network to generate relevant object proposals (Block-3), (ii) **proposal scoring:** the pooled object proposals (shown using indigo) are scored with sketch feature vector (shown using blue) to generate localization for the object of interest (Block-4). [Best viewed in color].

The proposed sketch-guided object localization is an end-to-end trainable framework which works in the following two stages: (i) query-guided object proposal generation (Section 3.2), (ii) proposal scoring (Section 3.3). Fig. 2 illustrates the overall architecture of the proposed framework.

3.2 Cross-modal Attention for Query-guided Object Proposal Generation

A popular object detection framework, viz. Faster R-CNN [24] uses a region proposal network (RPN) module to generate object proposals. We can use the same module for generating object proposals in our task. However, the RPN is not designed to utilize any information on object appearance available in the query images. Hence, the object proposals that are relevant to the sketch query may not even be generated, especially when the object of interest is of low resolution, occluded or hidden among objects that are better represented in the input images. Therefore, using a RPN in its vanilla form may not be useful in our pipeline. To solve the aforementioned problem, we propose *cross-modal attention* to incorporate the sketch information in the RPN to guide the proposal generation. Regions of interest (ROIs) are pooled from these region proposals

using a strategy similar to the Faster R-CNN, and a scoring function Θ is learned between these ROIs and sketch queries.

We will now describe our novel *cross-modal attention* scheme that we have introduced to generate object proposals relevant to an input query sketch. The attention module is trained to produce a weight map that provides high scores to the areas in an image that are similar to the given query sketch.

A sketch $s_c \in S$ of an object category $c \in C$ is used to query an image $i \in I$. To generate the feature representation of sketches and images, we use ResNet-50 models pretrained on Imagenet [4] and Quick-Draw [15] datasets for images and sketches respectively as a backbone. Suppose ϕ_I and ϕ_S represent these backbone models, then image and sketch feature maps are computed as:

$$i^{\phi_I} = \phi_I(i) \quad \text{and} \quad s_c^{\phi_S} = \phi_S(s_c), \quad (1)$$

where, $i^{\phi_I} \in \mathbb{R}^{w \times h \times d}$ and $s_c^{\phi_S} \in \mathbb{R}^{w' \times h' \times d}$ are the extracted image and sketch feature maps respectively. From these feature maps, the compatibility score is learned between the sketch and the image feature maps by first applying non-linear transformations as below:

$$i^{\psi_I} = \psi_I(i^{\phi_I}) \quad \text{and} \quad s_c^{\psi_S} = \psi_S(s_c^{\phi_S}). \quad (2)$$

A set of local feature vectors is formed by obtaining one vector at each location (m, n) in the image feature map i^{ψ_I} , where $m \in \{1, 2, \dots, w\}$ and $n \in \{1, 2, \dots, h\}$. Each vector represents a spatial region on the target image and the set gives us spatial distribution of the features. Subsequently, this is compared against a global representation of the sketch features. For image feature map i.e. $i^{\psi_I} \in \mathbb{R}^{w \times h \times d}$, the extracted set of feature vectors is represented as $L^i = \{\mathbf{L}_1^i, \mathbf{L}_2^i, \dots, \mathbf{L}_{w \times h}^i\}$ where $\mathbf{L}_j^i \in \mathbb{R}^{1 \times 1 \times d} \forall j \in \{1, 2, \dots, w \times h\}$.

In the case of sketches, a global representation of sketch feature maps is obtained via the global max pool (\mathcal{GMP}) operation, i.e.,

$$\mathbf{L}_g^{s_c} = \mathcal{GMP}(s_c^{\psi_S}), \quad (3)$$

where, $\mathbf{L}_g^{s_c} \in \mathbb{R}^{1 \times 1 \times d}$. A spatial compatibility score between $\mathbf{L}_j^i \in L^i$ and $\mathbf{L}_g^{s_c}$, is evaluated as follows:

$$\lambda(\mathbf{L}_j^i, \mathbf{L}_g^{s_c}) = \frac{(\mathbf{L}_j^i)^T (\mathbf{L}_g^{s_c})}{\mathcal{K}}, \quad (4)$$

where \mathcal{K} is a constant. For simplicity of notation, we will refer to the left hand side of eq.(4) as λ_{jg} from here onwards.

Note that these compatibility scores are generated as a spatial map, which can be understood as a weight map representing attention weights. Therefore, in order to obtain attended feature maps, we perform element-wise multiplication of these compatibility scores and the original image feature map at each spatial location, i.e.,

$$i_j^{a_I} = i_j^{\phi_I} \odot \lambda_{jg}, \forall j \in \{1, 2, \dots, w \times h\}. \quad (5)$$

This attention feature map aims to capture information about the location of objects in an image that shares high compatibility score with the sketch query.

Therefore, to incorporate this information, attention feature maps are concatenated along the depth with the original feature maps, i.e., $i_f^{\phi_I} = [(i^{a_I})^T; (i^{\phi_I})^T]^T$, where $i_f^{\phi_I} \in \mathbb{R}^{w \times h \times 2d}$. These concatenated feature maps are projected to a lower dimensional space to obtain the final feature maps, which are subsequently passed through the RPN to generate object proposals relevant to the sketch query.

3.3 Proposal Scoring

Once a small set of proposals, represented as R_i for $i \in I$, are pooled from all query-guided region proposals generated by RPN, a scoring function Θ is learned to rank these proposals with respect to the sketch query. To this end, during the training phase, each of these region proposals are labeled as foreground (1) if it has ≥ 0.5 intersection over union (IoU) with any of the ground-truth bounding boxes and the object in the ground-truth bounding box is the same class as sketch query, and background (0) otherwise. Then, we minimize a margin rank loss between the generated object proposals and the sketch query such that object proposals similar to the query sketch are ranked higher.

To learn the scoring function Θ , firstly, we generate feature vectors for the sketch query and the object proposals by taking a global mean pool operation on the sketch feature maps and object proposal feature maps, respectively. Each of the proposal feature vectors is concatenated with the sketch feature vector. These concatenated feature vectors are passed through a scoring function and the foreground probabilities of the proposals in context with the sketch query are predicted. Let a_k be the predicted foreground probability for proposal $r_k \in R_i$, and it is given by:

$$a_k = \Theta([g_m(r_k^{\phi_I})^T; g_m(s_c^{\phi_S})^T]^T), \quad (6)$$

where, $r_k^{\phi_I}$ is the feature map for $r_k \in R_i$ generated using standard Faster R-CNN protocols and $g_m : \mathbb{R}^{W \times H \times D} \rightarrow \mathbb{R}^{1 \times 1 \times D}$ is the global mean pool operation. Now, towards training the scoring function, a label $y_k = 1$ is assigned to r_k if it is part of any foreground bounding box and $y_k = 0$, otherwise. Motivated from [12] the loss function used in training is defined as:

$$L(R_i, s_c) = \sum_k \{y_k \max(m^+ - a_k, 0) + (1 - y_k) \max(a_k - m^-, 0) + L_{MR}^k\} \quad (7)$$

$$L_{MR}^k = \sum_{l=k+1} \{ \mathbb{1}_{[y_l=y_k]} \max(|a_k - a_l| - m^-, 0) + \mathbb{1}_{[y_l \neq y_k]} \max(m^+ - |a_k - a_l|, 0) \}, \quad (8)$$

where m^+ and m^- are positive and negative margins, respectively. The above loss function consists of two parts: (i) The first part of the loss (see eq.(7)) ensures that the proposals overlapping with the ground truth object locations are

predicted as foreground with high confidence. (ii) The second part of the loss function (see eq. (8)) is a margin-ranking loss that considers pairs of the proposals as input. It helps to further enforce a wider separation between foreground and background proposals in terms of the prediction probabilities, thereby improving the ranking of all the foreground proposals overlapping with the true location(s) of the object of interest. In addition to this loss, cross-entropy loss on the labeled (background or foreground) feature vectors of the region proposals and a regression loss on the predicted bounding box locations with respect to the ground truth bounding box are also used for training.

3.4 Multi-query Object Localization Setup

The quality of the sketches can hinder the object localization performance. In many cases, sketches can be abstract and may not often contain the structural attributes that differentiate one object from another. The noise present in such sketches, along with their abstract nature, makes the task of sketch-guided object localization extremely challenging. However, we observe that the sketches are diverse in quality and when produced by different creators, they often tend to capture complementary information on an object’s shape, attributes and appearance. This quality of complementarity can be leveraged towards improving the overall localization performance if information across these sketches is judiciously combined. With these motivations, we introduce the task of multi sketch query-guided object localization and present the following fusion strategies:

Feature fusion: We observed complementarity across filter responses for different input sketches containing the same object. Such complementarity can be leveraged by a suitable feature fusion strategy to obtain less noisy and more holistic representation of the sketched object. In this paper, we used the global max pool operation to fuse feature maps of different sketch queries. Let $\{s_c^1, \dots, s_c^N\}$ be the set of N sketches for the same object category $c \in C$ and the representation learned by the backbone network for the n^{th} sketch in this set is ${}^n s_c^{\phi_S}$. These feature maps are concatenated together to yield a composite feature map $R^{w \times h \times d \times N}$. A global max pool operation across all N channels is subsequently performed to obtain a fused feature map for all the N queries and the same is fed as input to the *cross-modal attention* model.

Attention Fusion: In an alternative strategy, to attenuate noise in the attention maps produced by individual sketches, we concatenate these maps as generated by multiple queries and subsequently perform a depth-wise average pooling operation to obtain the final attention map. The fused attention map, thus generated, is used as input to the object localization pipeline (Section 3.3).

4 Experiments and Results

4.1 Datasets

We use the following datasets for evaluating the performance of the proposed sketch-guided object localization framework.

Model	mAP	%AP@50
Modified Faster R-CNN	0.18	31.5
Matchnet [12]	0.28	48.5
Cross-modal attention (this work)	0.30	50.0

Table 1: **Results in one-shot common train-test categories setting on COCO *val2017* dataset.** Comparison of various object localization baseline methods with the proposed *cross-modal attention* model for sketch-guided object localization. We clearly outperform both the baselines based on mean average precision as well as % AP@50. For more details please refer to Section 4.4.

QuickDraw [15] is a large-scale sketch dataset containing 50 million sketches across 345 categories. Sketches drawn from a subset of this dataset (please refer to the next paragraphs) are used as query. QuickDraw sketches are stored as vector graphics and are rasterized before feeding into ResNet.

MS-COCO [19] is a large-scale image dataset which has been extensively used in object detection research. It has a total of 81 classes, including background class, with dense object bounding box annotations. There are a total of 56 classes which are common between MS-COCO and QuickDraw datasets. Therefore, in our object localization experiments, we randomly selected a total of 800K sketches across these common classes. We trained our model on COCO *train2017* and evaluated on COCO *val2017* dataset.

PASCAL VOC [6] is a popular object detection dataset with a total of 20 classes. Among these, nine classes common with the QuickDraw dataset are selected for our experiments. Our model is trained on the union of VOC2007 *train-val* and VOC2012 *train-val* sets and evaluated on VOC *test2007* set.

4.2 Baselines

We adapt the following popular models from the object detection and image-guided localization communities as baselines towards evaluating sketch-guided object localization performance and compare them with our proposed method.

Faster R-CNN [24] We adapt this object detector for query-guided object localization task. To this end, during training, we assign a 1 (or 0) class label to the object instance in the image if it belongs to the same (or different) class as the sketch query, and generate object proposals. Then, a binary classifier is used to classify each proposal as background or foreground. The query features are first concatenated with the region of interests (ROIs) features (pooled from region proposals) before passing it through the binary classifier. Additionally, we used a triplet loss to rank the object region proposal with respect to the sketch query. We refer to this baseline as *modified Faster R-CNN*.

Model	mAP
Modified Faster R-CNN	0.65
Matchnet [12]	0.61
Cross-modal attention (this work)	0.65

Table 2: **Results in one-shot common train-test categories setting on VOC *test2007* dataset.** Comparison of various object localization baseline methods with the proposed *Cross-modal attention* model for sketch-guided object localization. We clearly outperform the state-of-the-art object localization method [12] and the results are comparable to Faster R-CNN based baseline. Please refer to Section 4.4 for more details.

Matchnet [12] is a one-shot object localization method via image query. Here, non-local neural networks [34] and channel co-excitation [13] were used to incorporate the query information in the image feature maps. We adapt this recent method to directly work with a sketch query and treat it as a second baseline.

For both these baseline methods, the feature extractors for sketches and images are ResNet-50 models pretrained on Imagenet and QuickDraw respectively.

4.3 Experimental Setup

We used two ResNet-50 models pre-trained on Imagenet [4] and 5 million images of QuickDraw [15] to get the feature representation for images and sketches respectively. The images from MS-COCO and PASCAL VOC datasets are used as target images, and the sketches randomly-drawn from the common classes of QuickDraw are used as queries. We evaluate performance of our model under the following four settings: (i) one-shot common train-test categories, (ii) one-shot disjoint train-test categories, (iii) multi-query common train-test categories, and (iv) multi-query disjoint train-test categories.

Disjoint train-test experimental setting: Out of the 56 classes common across COCO and QuickDraw datasets, 42 and 14 classes are arbitrarily picked as ‘seen’ and ‘unseen’ categories, respectively. To ensure a one-shot disjoint train-test setting, the ‘seen’ and ‘unseen’ splits are mutually exclusive in terms of object classes as well as labeled bounding boxes present. Our model is trained solely on the ‘seen’ classes and only the ‘unseen’ classes are used for one-shot evaluation. Similarly for PASCAL VOC dataset, out of 9 classes common with QuickDraw, 3 and 6 are arbitrarily picked as ‘unseen’ and ‘seen’ categories respectively. The image encoder is pre-trained using Imagenet dataset excluding 14 ‘unseen’ as well as all related classes to these 14 classes obtained by matching their WordNet synsets. Similarly, the sketch encoder is pre-trained using all the QuickDraw classes except the 14 categories in the unseen set.

Model	Unseen Classes	Seen Classes
Modified Faster R-CNN	7.4	34.5
Matchnet [12]	12.4	49.1
Cross-modal attention (this work)	15.0	48.8

Table 3: **Results in one-shot disjoint train-test categories setting on MS-COCO *val2017* dataset.** We report % AP@50 scores in this table. Here, *unseen classes* contain the evaluation images and sketches of object categories that are not seen during training. Our proposed method significantly outperforms other baselines on unseen classes while achieving reasonably high performance for seen categories as well. Please refer to Section 4.4 for more details.

4.4 Results and Discussion

In this section, we report results on the four experimental settings discussed in Sec. 4.3. We then provide extensive discussions around merits and limitations of our proposed method in comparison to the chosen baselines.

One-shot common train-test categories: The results for the proposed method in this setting are shown in Table 1 for MS-COCO dataset. As shown, the proposed method outperforms both the baselines significantly. When compared to the modified Faster R-CNN baseline, our method performs significantly better. This is primarily because, unlike Faster R-CNN, our *cross-modal attention* framework effectively incorporates information from the query using spatial compatibility (attention) map to generate relevant region proposals. Moreover, our method also outperforms the Matchnet baseline because the non-local feature maps and channel co-excitation module in Matchnet are sensitive towards the domain gap present between query and image feature maps in our task. On the contrary, our *cross-modal attention* framework addresses this by explicitly computing a spatial compatibility map (attention map).

The results on the PASCAL VOC dataset are reported in Table 2. PASCAL VOC contains less number of training images ($\approx 9K$) with small variability across training classes in our setting (only 9 classes in the training set). The Faster R-CNN baseline is comparable to the proposed method indicating that sketch query-guided object localization is challenging without sufficient data. The Matchnet baseline model degrades in performance, indicating its inability to incorporate sketch information during proposal generation in the case of small data size and large domain gap.

One-shot disjoint train-test categories: In this setting, models are evaluated on the object categories which are not seen during training. The results are provided in Table 3 for MS-COCO dataset. We have selected model checkpoints that perform the best on *unseen* categories for each model. It is evident that one-shot object localization under disjoint train-test object categories is a hard problem as it leads to degradation in performance for all three models. Faster

Our models	Unseen (D)	Seen (D)	All (C)	mAP (C)
Cross-modal attention	15.0	48.8	50.0	0.30
+Query Fusion(5Q)	16.3	52.2	52.6	0.32
+Attention Fusion(5Q)	17.1	51.9	53.1	0.32

Table 4: **Results in multi-query common and disjoint train-test categories setting on MS-COCO *val2017* dataset.** Comparison of the proposed fusion strategies for 5 sketch queries. Here, 5Q, (C) and (D) represents 5 sketch queries, common train and test categories, and disjoint train and test categories respectively. % AP@50 is reported except the last column which reports mAP.

R-CNN baseline suffers most degradation because it does not incorporate sketch information during proposal generation and perform poorly in proposal generation for unseen categories. While Matchnet outperforms the modified Faster R-CNN baseline, it still performs significantly worse than the proposed method. Contrary to baseline methods, our model, by virtue of *cross-modal attention*, embeds query information in the image feature representation before generating region proposals. Therefore, it is inherently capable of generating relevant object proposals even when attempting to localize unseen object classes. As a result, our method reports superior performance on unseen categories while achieving reasonably high performance for seen categories as well.

Multi-query common and disjoint train-test categories: Here we discuss the effect of utilizing multiple sketch queries on the object localization performance. We evaluated our model using two fusion strategies presented in section 3.4 under the following two settings: (a) common train and test categories, (b) completely disjoint train and test object classes. The evaluation results of these fusion strategies are provided in Table 4. Both these fusion strategies provide consistent improvement in performance over the proposed cross-modal attention based baseline model, which just uses a single sketch query as input. This suggests that the fusion strategies are able to capture the complementary information present across the multiple sketches of the same object. Moreover, we also evaluated the proposed fusion strategies for sketch-guided object localization on *unseen* categories and observe that even in such cases, these fusion strategies are able to aid the proposed localization framework.

Utility of margin-based ranking loss: We analyze the effect of different components of loss function defined in eq.(7) and eq. (8) on the model’s performance using MS-COCO *val2017* dataset. The model which is trained with only foreground-background classification loss (first part of eq.(7)) gives %AP@50 of 40.0% as compared to 41.4%, achieved by the model trained on only the margin-ranking loss (L_{MR}) (eq.(8)). However, training the model on both these losses significantly improves the performance (50.0 %AP@50).




Fig. 3: Sketch queries and object localization results are shown along with the corresponding attention maps. The attention map produces high compatibility scores (shown in red color) for the regions on the natural image that contains object same as that in the sketch query. This leads to the localization (shown using a red bounding box on input image) of the relevant object. Please note that sketches have been enhanced for better visualization. **[Best viewed in color]**.

Visualizing the attention map and localization: In order to visualize the effect of sketch queries on the generated attention maps and the resulting localizations, we query an image with sketches of different object categories. As shown in Fig. 3, our novel *cross-modal attention* scheme is able to assign high compatibility scores to the ground-truth object locations in the natural image while at the same time ignoring instances of other object categories, thereby producing precise localizations.

5 Conclusion

In this paper, we have introduced a novel problem of sketch-guided object localization in natural images and presented an end-to-end trainable framework using novel *cross-modal attention* for this task. The effectiveness of the proposed attention scheme is substantiated by the significant performance improvement achieved over the baseline methods on two public benchmarks. We have extended the framework to a multi-query setup and proposed two fusion strategies towards the same. The superior performance of our framework over the baseline in this setting corroborates that the proposed fusion strategies are able to leverage the complementarity present across multiple sketches of the same object. We look forward to exciting future research in sketch-guided computer vision tasks inspired by the problem that we have introduced.

Acknowledgements

The authors would like to thank the Advanced Data Management Research Group, Corporate Technologies, Siemens Technology and Services Pvt. Ltd., and Pratiksha Trust, Bengaluru, India for partly supporting this research.

References

1. Cai, Z., Vasconcelos, N.: Cascade R-CNN: Delving into high quality object detection. In: CVPR (2018)
2. Caicedo, J.C., Lazebnik, S.: Active object localization with deep reinforcement learning. In: ICCV (2015)
3. Choe, J., Shim, H.: Attention-based dropout layer for weakly supervised object localization. In: CVPR (2019)
4. Deng, J., Dong, W., Socher, R., Li, L.J., Li, K., Fei-Fei, L.: Imagenet: A large-scale hierarchical image database. In: CVPR (2009)
5. Eitz, M., Hildebrand, K., Boubekeur, T., Alexa, M.: Sketch-based image retrieval: Benchmark and bag-of-features descriptors. *IEEE Transactions on Visualization and Computer Graphics* **17**(11), 1624–1636 (2010)
6. Everingham, M., Van Gool, L., Williams, C.K.I., Winn, J., Zisserman, A.: The pascal visual object classes (voc) challenge. *International Journal of Computer Vision* **88**(2), 303–338 (Jun 2010)
7. Girshick, R.: Fast R-CNN. In: ICCV (2015)
8. Girshick, R., Donahue, J., Darrell, T., Malik, J.: Rich feature hierarchies for accurate object detection and semantic segmentation. In: CVPR (2014)
9. Ha, D., Eck, D.: A neural representation of sketch drawings. arXiv preprint arXiv:1704.03477 (2017)
10. He, K., Gkioxari, G., Dollár, P., Girshick, R.: Mask R-CNN. In: ICCV (2017)
11. He, K., Zhang, X., Ren, S., Sun, J.: Spatial pyramid pooling in deep convolutional networks for visual recognition. *IEEE transactions on pattern analysis and machine intelligence* **37**(9), 1904–1916 (2015)
12. Hsieh, T.I., Lo, Y.C., Chen, H.T., Liu, T.L.: One-shot object detection with co-attention and co-excitation. In: NeurIPS (2019)
13. Hu, J., Shen, L., Sun, G.: Squeeze-and-excitation networks. In: CVPR (2018)
14. Hu, R., Collomosse, J.: A performance evaluation of gradient field hog descriptor for sketch based image retrieval. *Computer Vision and Image Understanding* **117**(7), 790–806 (2013)
15. Jongejan, J., Rowley, H., Kawashima, T., Kim, J., Fox-Gieg, N.: The quick, draw!-ai experiment, <https://quickdraw.withgoogle.com>, 2016.
16. Kong, T., Sun, F., Liu, H., Jiang, Y., Shi, J.: Foveabox: Beyond anchor-based object detector. arXiv preprint arXiv:1904.03797 (2019)
17. Li, B., Yan, J., Wu, W., Zhu, Z., Hu, X.: High performance visual tracking with siamese region proposal network. In: CVPR (2018)
18. Lin, T.Y., Goyal, P., Girshick, R., He, K., Dollár, P.: Focal loss for dense object detection. In: ICCV (2017)
19. Lin, T.Y., Maire, M., Belongie, S., Hays, J., Perona, P., Ramanan, D., Dollár, P., Zitnick, C.L.: Microsoft COCO: Common objects in context. In: ECCV (2014)
20. Liu, L., Shen, F., Shen, Y., Liu, X., Shao, L.: Deep sketch hashing: Fast free-hand sketch-based image retrieval. In: CVPR (2017)
21. Liu, W., Anguelov, D., Erhan, D., Szegedy, C., Reed, S., Fu, C.Y., Berg, A.C.: Ssd: Single shot multibox detector. In: ECCV (2016)
22. Redmon, J., Divvala, S., Girshick, R., Farhadi, A.: You only look once: Unified, real-time object detection. In: CVPR (2016)
23. Redmon, J., Farhadi, A.: Yolov3: An incremental improvement. arXiv preprint arXiv:1804.02767 (2018)

24. Ren, S., He, K., Girshick, R., Sun, J.: Faster R-CNN: Towards real-time object detection with region proposal networks. In: *NeurIPS (2015)*
25. Sangkloy, P., Burnell, N., Ham, C., Hays, J.: The sketchy database: learning to retrieve badly drawn bunnies. *ACM Transactions on Graphics (TOG)* **35**(4), 1–12 (2016)
26. Sermanet, P., Eigen, D., Zhang, X., Mathieu, M., Fergus, R., LeCun, Y.: Overfeat: Integrated recognition, localization and detection using convolutional networks. *arXiv preprint arXiv:1312.6229* (2013)
27. Sharif Razavian, A., Azizpour, H., Sullivan, J., Carlsson, S.: Cnn features off-the-shelf: an astounding baseline for recognition. In: *CVPRW (2014)*
28. Shen, Y., Liu, L., Shen, F., Shao, L.: Zero-shot sketch-image hashing. In: *CVPR (2018)*
29. Sivic, J., Zisserman, A.: Video Google: A text retrieval approach to object matching in videos. In: *ICCV (2003)*
30. Song, J., Yu, Q., Song, Y.Z., Xiang, T., Hospedales, T.M.: Deep spatial-semantic attention for fine-grained sketch-based image retrieval. In: *ICCV (2017)*
31. Teh, E.W., Rochan, M., Wang, Y.: Attention networks for weakly supervised object localization. In: *BMVC (2016)*
32. Uijlings, J.R., Van De Sande, K.E., Gevers, T., Smeulders, A.W.: Selective search for object recognition. *International Journal of Computer Vision* **104**(2), 154–171 (2013)
33. Wan, J., Wang, D., Hoi, S.C.H., Wu, P., Zhu, J., Zhang, Y., Li, J.: Deep learning for content-based image retrieval: A comprehensive study. In: *ACM MM (2014)*
34. Wang, X., Girshick, R., Gupta, A., He, K.: Non-local neural networks. In: *CVPR (2018)*
35. Xu, P., Huang, Y., Yuan, T., Pang, K., Song, Y.Z., Xiang, T., Hospedales, T.M., Ma, Z., Guo, J.: Sketchmate: Deep hashing for million-scale human sketch retrieval. In: *CVPR (2018)*
36. Xu, P., Joshi, C.K., Bresson, X.: Multi-graph transformer for free-hand sketch recognition. *arXiv preprint arXiv:1912.11258* (2019)
37. Yan, X., Chen, Z., Xu, A., Wang, X., Liang, X., Lin, L.: Meta R-CNN: Towards general solver for instance-level low-shot learning. In: *ICCV (2019)*
38. Yu, Q., Liu, F., Song, Y.Z., Xiang, T., Hospedales, T.M., Loy, C.C.: Sketch me that shoe. In: *CVPR (2016)*
39. Yu, Q., Yang, Y., Liu, F., Song, Y.Z., Xiang, T., Hospedales, T.M.: Sketch-a-net: A deep neural network that beats humans. *International Journal of Computer Vision* **122**(3), 411–425 (2017)
40. Zhang, J., Shen, F., Liu, L., Zhu, F., Yu, M., Shao, L., Tao Shen, H., Van Gool, L.: Generative domain-migration hashing for sketch-to-image retrieval. In: *ECCV (2018)*