

Common Object Discovery as Local Search for Maximum Weight Cliques in a Global Object Similarity Graph

Cong Rao¹, Yi Fan^{2,3}, Kaile Su³, and Longin Jan Latecki¹

¹ Temple University, Philadelphia, USA
{cong.rao,latecki}@temple.edu

² Guangxi Key Lab of Trusted Software,
Guilin University of Electronic Technology, Guilin 541004, China

³ Griffith University, Brisbane, Australia
{yi.fan4,k.su}@griffithuni.edu.au

Abstract. In this paper, we consider the task of discovering the common objects in a set of images. Initially, object candidates are generated in each image and an undirected weighted graph is constructed over all the candidates. Each candidate serves as a node in the graph while the weight of the edge describes the similarity between the corresponding pair of candidates. The problem is then expressed as a search for the Maximum Weight Clique (MWC) in this graph. The MWC corresponds to a set of object candidates sharing maximal mutual similarity, and each node in the MWC represents a discovered common object class across the images. Since the problem of finding MWCs is NP-hard, the research of the MWC problem focuses on developing various heuristics for finding good cliques within a reasonable time limit. We utilize a recently very popular class of heuristics called local search methods. They search for MWCs directly in the discrete domain of the solution space. The proposed approach is evaluated on the PASCAL VOC image dataset and the YouTube-Objects video dataset, and it demonstrates superior performance over recent state-of-the-art approaches.

Keywords: Common Object Discovery · Visual Similarity · Maximum Weight Clique · Local Search Algorithm

1 Introduction

We focus on the task of common object discovery, which aims at discovering the objects of the same class in an image collection. Co-localizing objects in unconstrained environment is challenging. For images in the real-world applications, such as those in the PASCAL datasets [8, 9], the objects of the same class may look very different due to viewpoint, occlusion, deformation, illumination, *etc.* Also, there could be considerable diversities within certain object class such as human beings, for their differences in gender, age, costume, hair style or even skin color. Besides, there could be multiple common objects in the given set of

images, thus the definition of “common” may be ambiguous. For example, for the first two images in Fig. 1, the common object class could be either human or bicycle. In addition, the efficiency of the involved method is very significant in time sensitive applications such as object co-localization in large collections of images or video streams.



Fig. 1. Given a set of object candidates generated from an image collection (left), our goal is to find common objects by searching for the maximum weight clique in the associated graph. Each node in the clique (right) corresponds to a discovered common object.

To achieve robust and efficient object co-localization, we formulate the task as a Maximum Weight Clique (MWC) problem. It aims at finding a group of objects that are most similar to each other, which corresponds to a MWC in the associated graph. The nodes in the graph correspond to the object candidates generated from the given image collection, while the weight on an edge indicates how similar two given candidates are. We can discover a set of common objects by finding the MWC in the associated graph. Each node in the MWC is a discovered common object across the images. Figure 1 illustrates the idea.

For an undirected weighted graph $G = (V, E)$, where V is the set of vertices and E is the set of edges, a clique C is a subset of vertices in V and each pair of vertices in C is connected by an edge in E . The Maximum Weight Clique (MWC) problem is to find a clique C which maximizes

$$w(C) = \sum_{v_i \in C} w_V(v_i) + \sum_{v_i, v_j \in C} w_E(v_i, v_j), \quad (1)$$

where $w_V : V \rightarrow \mathbb{R}$ and $w_E : E \rightarrow \mathbb{R}$ are the weight functions for the vertices and edges respectively. Successfully solving the MWC problem has various applications not only in computer vision [17, 12] but in many different domains from wireless to social networks [21, 19].

The main contributions of our paper include:

1. We address the task of object co-localization as a well-defined MWC problem in the associated graph. It provides a practical and general solution for research and applications related to the MWC problem.
2. We develop a hashing based mechanism to detect the revisiting of the local optimum in the local search based MWC solver [36]. It can alleviate the cycling issue in the optimization process.

3. The Region Proposal Network (RPN) [24] is applied for efficiently generating the object candidates. The candidates are then re-ranked to improve the robustness against the background noise.
4. A Triplet Network (TN) is learned to obtain the feature embeddings of the object candidates, so as to construct a reliable affinity measure between the candidates.
5. The performance is evaluated on the PASCAL VOC 2007 image dataset [8] and the YouTube-Objects video dataset [13]. Superior performance is obtained compared to recent state-of-the-art methods.

2 Related Works

The problem of common object discovery has been investigated extensively in the past few years. Papazoglou *et al* [20] view the task as a foreground object mining problem, where optical flow is used to estimate object motion and the Gaussian mixture model is utilized to capture the appearance of the foreground and background. Cho *et al* [6] tackle the problem using a part-based region matching method, where a probabilistic Hough transform is used to evaluate the quality of each candidate correspondence. Joulin *et al* [12] extend the method in [6] to co-localize objects in video frames, and a Frank-Wolfe algorithm is used to optimize the proposed quadratic programming problem. In [35], Edge Box [39] is employed to generate the object candidates. The common object discovery problem is then addressed under the Multiple Instance Learning (MIL) framework. Rochan *et al* [26] present a method for transferring existing appearance models to unseen objects and apply it to localize common objects in images and videos. Zhang *et al* [38] apply a part-based object detector and a motion aware region detector to generate object candidates. The problem is then formulated as a joint assignment problem and the solution is refined by inferring shape likelihoods afterwards. Bilen *et al* [3] extract deep neural network features upon Selective Search [31] windows and use Latent SVM (LSVM) to discover the common objects. Kwak *et al* [14] also focus on the problem of localizing dominant objects in videos, where an iterative process of detection and tracking is applied. Li *et al* [15] devise an entropy-based objective function to learn a common object detector, and they address the task with a Conditional Random Field (CRF) model. Wang *et al* [34] focus on segmenting common objects in videos, and they aim at minimizing the cost of assigning labels to the super-pixels in each frame. Wei *et al* [37] perform Principal Component Analysis (PCA) on the convolutional feature maps of all the images, and locate the most correlated regions across the images. Wang *et al* [32] use segmentations produced by Fully Convolutional Networks (FCN) as object candidates. Then they discover common objects by solving a N -Partite Graph Matching problem.

Many of these methods explicitly or implicitly employ graph based models to interpret the task of object co-localization. Compared to these methods, we apply different problem formulations, focus on different visual primitives, utilize different appearance models, and adopt different metrics to evaluate the consistency of the common objects. In this paper, an undirected weighted graph is

first constructed over the given set of images, modeling the visual affinities between the object candidates. We find common objects as the Maximum Weight Clique (MWC) in this graph, where each node in the clique corresponds to a detected common object across the images. The MWC problem is NP-hard and it is difficult to obtain a global optimal solution, thus often relaxation methods are applied so that problem can be solved in the continuous domain [17]. In this paper, we focus on the MWC solvers in the discrete domain to tackle the problem. Generally, there are two types of algorithms that work directly in the discrete domain to solve the MWC problem: methods based on global optimization [2, 11] and those based on local search [23, 18, 36]. While the global optimization methods guarantee an optimal solution, they do not scale up to real problems on large graphs. In contrast, the local search methods often yield solutions close to global optimum on real graphs.

3 Problem formulation

Given a set of N images $\mathcal{I} = \{I_1, I_2, \dots, I_N\}$, we generate a set of possible object candidates from all images $\mathcal{B} = \{\mathbf{b} \mid \mathbf{b} \in \mathcal{P}(I), I \in \mathcal{I}\}$, where $\mathcal{P}(I)$ is the set of object candidates extracted from image I and \mathbf{b} is a bounding box of that object candidate. Suppose n_i object candidates are extracted from image I_i , then a total number of $|\mathcal{B}| = \sum_{i=1}^N n_i$ candidates will be generated from the image collection \mathcal{I} . We denote $n = |\mathcal{B}|$ in the remainder of the paper. Let $o(\mathbf{b}_i)$ be the probability of some bounding box \mathbf{b}_i containing the common object, and let $s(\mathbf{b}_i, \mathbf{b}_j)$ represent the similarity between two object candidates in \mathbf{b}_i and \mathbf{b}_j , then the task of object co-localization can be formulated as finding an optimal subset $\mathcal{B}^* \subset \mathcal{B}$ such that

$$w(\mathcal{B}^*) = \sum_{\mathbf{b}_i \in \mathcal{B}^*} o(\mathbf{b}_i) + \sum_{\mathbf{b}_i, \mathbf{b}_j \in \mathcal{B}^*, \mathbf{b}_i \neq \mathbf{b}_j} s(\mathbf{b}_i, \mathbf{b}_j) \quad (2)$$

is maximized, with the constraint that at most one object candidate can be selected from each image.

Further, we assign a label $x_i \in \{0, 1\}$ to each object candidate \mathbf{b}_i , and $x_i = 1$ means that the object candidate \mathbf{b}_i is selected in the subset \mathcal{B}^* . Thus, an indicator vector $\mathbf{x} \in \{0, 1\}^n$ is used to identify the common objects discovered in \mathcal{B} . Besides, an affinity matrix $A \in \mathbb{R}^{n \times n}$ is constructed, where

$$A_{ii} = o(\mathbf{b}_i), \forall \mathbf{b}_i \in \mathcal{B}, \quad \text{and} \quad A_{ij} = s(\mathbf{b}_i, \mathbf{b}_j), \forall \mathbf{b}_i, \mathbf{b}_j \in \mathcal{B}. \quad (3)$$

Here we assume the similarity metric $s(\mathbf{b}_i, \mathbf{b}_j)$ is symmetric and non-negative, namely $A_{ij} = A_{ji} \geq 0$. On the other hand, we remove the edge between object candidates \mathbf{b}_i and \mathbf{b}_j if they are present in the same image, hence they cannot be simultaneously selected in \mathcal{B}^* . Then the problem in (2) can be expressed as finding an optimal indicator vector $\mathbf{x} \in \{0, 1\}^n$, such that $\mathbf{x}^T A \mathbf{x}$ is maximized. Hence the selected nodes in \mathcal{B}^* correspond to a MWC in the constructed graph, and they represent the discovered set of common objects. To summarize, the

overall objective function of the MWC problem can be written in the matrix form as

$$\mathbf{x}^* = \underset{\mathbf{x}}{\operatorname{argmax}} \mathbf{x}^T A \mathbf{x}, \quad \text{s. t. } \mathbf{x} \in \{0, 1\}^n. \quad (4)$$

To this end, the task of object co-localization is formulated as a Maximum Weight Clique (MWC) problem as described in (1).

4 Graph Construction

4.1 Object candidates generation

The nodes in the associated graph correspond to the object candidates in all the images. We expect those candidates to cover as many foreground objects as possible. Meanwhile, the total number of candidates will also influence the search space for the MWC. Therefore, our first priority is to find a proper method to extract the object candidates. The Region Proposal Networks (RPN) [24] is used in our approach to generate rectangular object candidates from each image. We use the raw RPN proposals in the intermediate stage and apply Non-Maximum Suppression (NMS) [27] to remove redundant boxes. We choose the top- K scoring proposals from each image to construct the associated graph for computational efficiency. We consider two different proposal scoring measures. The first one is commonly used and is based on RPN objectness score of each object candidate. RPN also generates a vector of class likelihoods for each object candidate, we propose to re-rank the object candidates according to the entropy of the class distribution. The entropy is low only if the class distribution is very different from uniform, namely when just a few class probabilities are high while the others are low. Thus small entropy values in this setting serve a similar purpose as the objectness score but tend to be more accurate. Hence we can re-rank the raw RPN proposals according to the entropy, and select the top- K scoring boxes as object candidates in each image.

4.2 Common object likelihood

For object co-localization, the underlying class of the common object is not known in advance. Thus, the object prior $o(\mathbf{b})$ is difficult to estimate. A possible way is to set the object prior as the objectness score of \mathbf{b} . But this can be problematic when \mathbf{b} contains an object but not the common one. Thus, it may lead to unexpected results if the objectness score is directly used, as observed in [30]. Therefore, we set the contribution of the unary term to the objective function (2) to zero, *i.e.*,

$$A_{ii} = o(\mathbf{b}_i) = 0, \forall \mathbf{b}_i \in \mathcal{B}.$$

In the case of object co-localization, this means we focus on an undirected graph with edge weight only. However, our formulation and MWC problem solver can be applied to other tasks where both vertex weights and edge weights are present.

4.3 Object representation and similarity

The edge weights in the associated graph represent visual similarity between the selected object candidates. Thus, we need an accurate way to represent the object candidates and evaluate their similarities. In this paper, we employ the Triplet Network framework to learn deep feature embeddings of the object candidates. Suppose a pre-trained Convolutional Neural Network (CNN) is selected to extract the deep features $f(\mathbf{b}; \mathbf{w})$ for each object candidate $\mathbf{b} \in \mathcal{B}$, where \mathbf{w} is the set of parameters of the CNN. In this framework, a set of triplets is then constructed for fine-tuning the parameters \mathbf{w} . Each triplet consists of a reference object \mathbf{b}_r , a positive object \mathbf{b}_p and a negative object \mathbf{b}_n . Namely, \mathbf{b}_r and \mathbf{b}_p represent a pair of similar objects, while \mathbf{b}_r and \mathbf{b}_n are a pair of dissimilar objects. Two objects are viewed as similar if they belong to the same category and otherwise dissimilar. Then, the hinge loss of a triplet is defined as

$$l(\mathbf{b}_r, \mathbf{b}_p, \mathbf{b}_n) = \max\{0, \lambda + s(\mathbf{b}_r, \mathbf{b}_n) - s(\mathbf{b}_r, \mathbf{b}_p)\}, \quad (5)$$

where λ is a margin threshold controlling how different $s(\mathbf{b}_r, \mathbf{b}_n)$ and $s(\mathbf{b}_r, \mathbf{b}_p)$ should be. The goal of the Triplet Network learning is to find a set of optimal parameters \mathbf{w} , such that the sum of the hinge loss of all triplets

$$L(\mathcal{T}) = \sum_{(\mathbf{b}_r, \mathbf{b}_p, \mathbf{b}_n) \in \mathcal{T}} l(\mathbf{b}_r, \mathbf{b}_p, \mathbf{b}_n) \quad (6)$$

is minimized over a training set of triplets \mathcal{T} . Namely, in the specified metric space, the learning process makes similar objects closer to each other, while dissimilar objects are pushed away. In the triplet hinge loss $l(\mathbf{b}_r, \mathbf{b}_p, \mathbf{b}_n)$, frequently used similarity metrics include dot-product (the linear kernel) and the Euclidean distance. But the output ranges of these metrics are not bounded, and this may invalidate the margin threshold λ in the loss function, as observed in [5]. In addition, more complex metrics can be also used here, such as the polynomial kernel and the Gaussian kernel (the RBF kernel). But there are a few more parameters in these kernel functions and they have to be chosen wisely. For simplicity, we define $s(\mathbf{b}_i, \mathbf{b}_j)$ as the cosine similarity between two CNN feature vectors $f(\mathbf{b}_i; \mathbf{w})$ and $f(\mathbf{b}_j; \mathbf{w})$, namely

$$A_{ij} = s(\mathbf{b}_i, \mathbf{b}_j) = \frac{f(\mathbf{b}_i; \mathbf{w})^T f(\mathbf{b}_j; \mathbf{w})}{\|f(\mathbf{b}_i; \mathbf{w})\| \|f(\mathbf{b}_j; \mathbf{w})\|}, \quad (7)$$

since the metric is already neatly bounded and parameter free. The parameters \mathbf{w} in the overall loss function (6) can be updated via the standard Stochastic Gradient Descent (SGD) method. The goal of this process is to learn a feature embedding such that similar objects are closer to each other in the metric space while dissimilar objects are pushed away.

5 The MWC problem solver

We use a local search based method to solve the MWC problem (4). The local search usually moves from one clique to another until it reaches the cutoff, then

the best clique found in the procedure is kept as the solution. The pipeline of our MWC solver is summarized in Algorithm 1.

Compared to RSL and RRWL in [10], our algorithm starts from a random single-vertex clique, while they generate a random maximal clique. This is particularly useful when the run-time is restricted. Besides, while RSL and RRWL restart when a solution is revisited in the so-called *first growing step*, our algorithm simply restarts when a local optimum is revisited. In this way, the solver will spend less time on searching the local area that has been visited intensively.

5.1 Detecting Revisiting via A Hash Function

In recent methods, the local search typically moves in a deterministic way, *i.e.*, no randomness exists in this process. Thus, a sequence of steps from a previously visited local optimum would be simply repeated, and it may not improve the best clique found so far. Hence, we improve this kind of methods by introducing a cycle elimination based restart strategy, where a hash table is used to approximately detect the revisiting of a local optimum. Given a candidate solution \mathcal{B}_c^* and a prime number p , we define the hash value of \mathcal{B}_c^* as

$$\text{hash}(\mathcal{B}_c^*) = \left(\sum_{\mathbf{b}_i \in \mathcal{B}_c^*} 2^i \right) \bmod p, \quad (8)$$

where $i \in \{1, 2, \dots, n\}$ is the index of \mathbf{b}_i in the entire object candidate set \mathcal{B} . If p is large enough, the chance of collision is negligible. The parameter p can be set according to the memory capacity of the machine. In the proposed algorithm, the revisiting of a local optimum is detected by checking whether the respective hash entry has been visited. If the local optimum was not visited before, the local search continues. Otherwise, the solver will be restarted and try to look for a better solution.

5.2 Scoring Functions and Candidate Nodes

Given an undirected weighted graph $G = (V, E)$, we describe our approach to finding the MWC in Algorithm 1. To begin with, we first introduce some notations used in our algorithm. In the local search for the MWC, the **add** operation adds a new node to the current clique C . The **drop** operation drops an existing node from the current clique C . The **swap** operation swaps two nodes from inside and outside the current clique C . Each operation returns a new clique as the current solution, which maximizes the gain of the clique weight. Suppose $w(C)$ is the weight of a clique C defined in Eq. (1), then for the **add** and **drop** operation, the gain of adding and dropping a node v is computed as

$$\text{score}(v, C) = \begin{cases} w(C \cup \{v\}) - w(C) & \text{if } v \notin C; \\ w(C \setminus \{v\}) - w(C) & \text{if } v \in C. \end{cases} \quad (9)$$

For **swap** operation, the gain of clique weight when swapping two nodes (u, v) is

$$\text{score}(u, v, C) = w(C \setminus \{u\} \cup \{v\}) - w(C), u \in C, v \notin C, (u, v) \notin E. \quad (10)$$

We denote the set of neighbors of a vertex v as $\mathcal{N}(v) = \{u | (u, v) \in E\}$. To ensure that the local search always maintains a clique, we define two operand sets as below. Firstly for a clique C , we define the set of candidate nodes for the **add** operation as

$$S_{add}(C) = \begin{cases} \{v | v \notin C, v \in \mathcal{N}(u), \forall u \in C\} & \text{if } |C| > 0; \\ \emptyset, & \text{otherwise.} \end{cases}$$

Secondly, the set of candidate node pairs for the **swap** operation is defined as

$$S_{swap}(C) = \begin{cases} \{(u, v) | u \in C, v \notin C, (u, v) \notin E, v \in \mathcal{N}(w), \forall w \in C \setminus \{u\}\} & \text{if } |C| > 1; \\ \emptyset, & \text{otherwise.} \end{cases}$$

To maximize the gain of clique weight in each step, the **add** operation adds a node v^* to the current clique C such that

$$v^* = \underset{v}{\operatorname{argmax}} \operatorname{score}(v, C), v \in S_{add}(C).$$

The **drop** operation drops a node

$$v^* = \underset{v}{\operatorname{argmax}} \operatorname{score}(v, C), v \in C,$$

from the current clique C . The **swap** operation swaps two nodes (u^*, v^*) such that

$$(u^*, v^*) = \underset{(u, v)}{\operatorname{argmax}} \operatorname{score}(u, v, C), (u, v) \in S_{swap}(C).$$

5.3 The Strong Configuration Checking Strategy

On the other hand, we apply the Strong Configuration Checking (SCC) strategy [36] to avoid revisiting a solution too early. The main idea of the SCC strategy works as follows. After a vertex v is dropped or swapped from a clique C , it can be added or swapped back into C only if one of its neighbors is added into C . Suppose $\operatorname{confChange}(v)$ is an indicator function of node v , where $\operatorname{confChange}(v) = 1$ means v is allowed to be added or swapped into the candidate solution and $\operatorname{confChange}(v) = 0$ means v is forbidden to be added or swapped into the candidate solution, then the SCC strategy specifies the following rules:

1. Initially $\operatorname{confChange}(v)$ is set to 1 for each vertex v ;
2. When v is added, $\operatorname{confChange}(u)$ is set to 1 for all $u \in \mathcal{N}(v)$;
3. When v is dropped, $\operatorname{confChange}(v)$ is set to 0;
4. When $(u, v) \in S_{swap}(C)$ are swapped, $\operatorname{confChange}(u)$ is set to 0.

Algorithm 1: Our MWC Problem Solver

Input : An undirected weighted graph $G = (V, E)$ and a time limit t
Output: A clique C^* with the maximum clique weight

- 1 $C^* \leftarrow C \leftarrow \emptyset$; $lastStepImproved \leftarrow true$;
- 2 $step \leftarrow 1$; $confChange(v) \leftarrow 1, \forall v \in V$;
- 3 **while** *elapsed time* $< t$ **do**
- 4 **if** $C = \emptyset$ **then** add a random vertex into C ;
- 5 $v \leftarrow \operatorname{argmax}_v score(v, C), v \in S_{add}(C)$, s.t. $confChange(v) = 1$;
- 6 $(u, u') \leftarrow \operatorname{argmax}_{(u, u')} score(u, u', C), (u, u') \in S_{swap}(C)$, s.t.
 $confChange(u') = 1$;
- 7 **if** $v \neq null$ **then**
- 8 **if** $(u, u') = (null, null)$ or $score(v) > score(u, u')$ **then**
- 9 $C \leftarrow C \cup \{v\}$;
- 10 **else**
- 11 $C \leftarrow C \setminus \{u\} \cup \{u'\}$;
- 12 $lastStepImproved \leftarrow true$;
- 13 **else**
- 14 **if** $(u, u') = (null, null)$ or $score(u, u') < 0$ **then**
- 15 **if** $lastStepImproved = true$ **then**
- 16 **if** $w(C) > w(C^*)$ **then** $C^* \leftarrow C$;
- 17 **if** *hash(C) is already marked* **then**
- 18 Drop all the vertices in C ;
- 19 **continue**;
- 20 Label *hash(C)* as marked;
- 21 $lastStepImproved \leftarrow false$;
- 22 **else**
- 23 $lastStepImproved \leftarrow true$;
- 24 $v' \leftarrow \operatorname{argmax}_{v'} score(v', C), v' \in C$;
- 25 **if** $(u, u') = (null, null)$ or $score(v', C) > score(u, u', C)$ **then**
- 26 $C \leftarrow C \setminus \{v'\}$;
- 27 **else**
- 28 $C \leftarrow C \setminus \{u\} \cup \{u'\}$;
- 29 Apply the Strong Configuration Checking (SCC) strategy;
- 30 $step++$;
- 31 **return** C^* ;

6 Experiments

To evaluate the performance of our method in comparison to other approaches, experiments are conducted on the PASCAL VOC 2007 image dataset [8] and the YouTube-Objects video dataset [13]. The standard PASCAL criterion Intersection over Union (IoU) is adopted for evaluation. Namely, a predicted bounding

box \mathbf{b}^p is correct if $IoU(\mathbf{b}^p, \mathbf{b}^{gt}) = \frac{area(\mathbf{b}^p \cap \mathbf{b}^{gt})}{area(\mathbf{b}^p \cup \mathbf{b}^{gt})} > 0.5$, where \mathbf{b}^{gt} is a ground-truth annotation of the bounding box. Finally, the percentage of images with correct object localization (CorLoc) [15] is used as the evaluation protocol. Our method is denoted as pMWC for the proposed MWC solver.

6.1 Implementation details

Our experiments are carried out on a desktop machine with two Intel(R) Core(TM) i7 CPUs (2.80GHz) and 64 GB memory. A GeForce GTX Titan X GPU is used for training and testing related deep neural networks. The proposed MWC solver is implemented in C/C++. The deep learning framework Caffe is utilized as a carrier for building the Region Proposal Network and the Triplet Network. The pipeline of the system is organized in MATLAB with some utilities written as MEX files, due to the efficiency for high level data management and visualization. The default parameters are used to learn RPN and generate the object candidates. A threshold of 0.5 is used for the NMS process to remove redundant object proposals. The best $K = 20$ object candidates are selected in each image. We set $\lambda = 0.25$ in the hinge loss (5) of a triplet. The prime number p in the hash function (8) is set to $10^9 + 7$, thus the hash table consumes around 1 GB memory. The RPN and Triplet Network in our method are built upon the VGG-f model [29] as well as the VGG-16 model [4]. Compared to the VGG-16 model, the structure of the VGG-f model is much simpler thus more computationally efficient. The VGG-f and VGG-16 models are pre-trained on the ImageNet dataset [28] and fine-tuned on the Microsoft COCO dataset [16]. All parameters are fixed the same in the experiments unless explicitly stated otherwise.

6.2 Experiments on the PASCAL07 dataset

The PASCAL VOC 2007 dataset [8] is used to evaluate the performance of object co-localization in images. The dataset is split as a training-validation set and a test set, each with about 5,000 images in 20 classes. We follow [15] to construct a test set from the training-validation set and denote it as PASCAL07. This is fine in our framework, since our RPN and Triplet networks are not trained on this dataset but on ImageNet and COCO datasets as stated in Sec. 6.1.

We first compare the co-localization accuracy of different MWC problem solvers on the PASCAL07 dataset in Table 1. The graph instances of these MWC problems are constructed based on the VGG-16 model. Since the PASCAL07 dataset has images from 20 different classes, we construct 20 different graphs, one graph for each image class. For the experiments on the PASCAL07 dataset, the average number of nodes in the constructed graphs is 6081.67, and the average number of edges is 2.16×10^7 . The average density of the graphs is 0.9962. Different solvers are evaluated on exactly the same MWC problem instances constructed by our co-localization framework. As randomized processes may exist in different methods, the reported accuracy is taken as the average over 10 runs with different seeds for the random number generator.

Method	Aero	Bike	Bird	Boat	Bottle	Bus	Car	Cat	Chair	Cow	Table	Dog	Horse	Motor	Person	Plant	Sheep	Sofa	Train	TV	Avg
MC [17]	63.0	45.7	56.1	51.9	14.3	47.8	71.9	61.1	36.2	72.3	46.0	57.2	61.3	79.6	62.3	34.3	69.8	39.3	59.4	9.8	52.0
LSCC [36]	61.3	68.3	61.2	48.1	16.8	67.7	76.9	58.5	41.8	72.3	24.5	63.9	68.3	75.5	69.0	28.6	76.0	47.2	62.1	68.4	57.8
[36] + BMS	63.9	68.3	60.9	50.3	49.2	66.7	76.9	59.3	41.1	72.3	23.0	65.1	68.3	77.1	69.8	27.8	76.0	45.9	62.8	68.4	59.7
TBMA [1]	62.2	69.5	62.1	52.5	18.4	71.5	78.5	61.1	49.2	70.9	30.0	62.0	69.7	80.8	66.0	49.4	70.8	50.2	63.6	68.0	60.3
pMWC	64.7	58.4	60.3	54.1	52.0	71.0	79.2	63.8	43.1	71.6	40.5	64.4	72.1	84.9	69.5	45.3	75.0	51.1	66.3	71.1	62.9

Table 1. Co-localization CorLoc (%) of different MWC solvers on the PASCAL07 dataset.

CNN Backbone	RPN+Objectness	RPN+entropy	Triplet Loss Fine-tuning
Pre-trained VGG-f Model	31.2	53.8	59.3
Pre-trained VGG-16 Model	33.1	56.1	62.9

Table 2. Co-localization CorLoc (%) of different strategies on the PASCAL07 dataset.

For the method [17], it solves the MWC problem in the relaxed continuous domain and a modified Frank-Wolfe algorithm is proposed to attack the problem. Similar to our approach, the solver TBMA [1] also solves the MWC problem directly in the discrete domain. Compared to their solver, our solver will restart if a local optimum is revisited, while TBMA will restart if the solution quality has not been improved for a specified number of steps. As for the solver LSCC [36], originally it is dedicated to solve the MWC problems where the edge weights are absent. Namely, the add, swap or drop operations change the weight of a clique considering related vertex weights only. Here we modify it so that the edge weights are taken into account in these operations. Two versions of the LSCC solver in the original paper are evaluated, and they serve as the baseline results in our experiment. The experiments justify our choice of the MWC problem solver, which improves the accuracy of object co-localization.

The co-localization accuracy of different object candidate generation and feature embedding methods on the PASCAL07 dataset are compared in Table 2. Different CNN models are used to extract the object candidate features, then the cosine similarity is applied on these deep neural network features. It shows that the proposed re-ranking the object candidates in each image according to the entropy of class distribution leads to significantly better results than the commonly used re-ranking based on objectness. With the involvement of the Triplet Network learning framework, the co-localization performance improves further. The experiments validate that the performance of the object co-localization is benefited from the proper choice of the object candidate generation and feature embedding scheme.

The co-localization accuracies of different object co-localization methods on the PASCAL07 dataset are reported in Table 3. The results of the compared methods are directly taken from the corresponding literatures. Among these methods using deep CNN features as visual descriptors [15, 35, 3, 25, 37, 7, 26], our method demonstrates superior results over recent state-of-the-art methods. The experiments confirm the effectiveness of the proposed object co-localization framework.

Method	Aero	Bike	Bird	Boat	Bottle	Bus	Car	Cat	Chair	Cow	Table	Dog	Horse	Motor	Person	Plant	Sheep	Sofa	Train	TV	Avg
Joulin <i>et al</i> [12]	32.8	17.3	20.9	18.2	4.5	26.9	32.7	41.0	5.8	29.1	34.5	31.6	26.1	40.4	17.9	11.8	25.0	27.5	35.6	12.1	24.6
Cho <i>et al</i> [6]	50.3	42.8	30.0	18.5	4.0	62.3	64.5	42.5	8.6	49.0	12.2	44.0	64.1	57.2	15.3	9.4	30.9	34.0	61.6	31.5	36.6
Li <i>et al</i> [15]	73.1	45.0	43.4	27.7	6.8	53.3	58.3	45.0	6.2	48.0	14.3	47.3	69.4	66.8	24.3	12.8	51.5	25.5	65.2	16.8	40.0
Wang <i>et al</i> [35]	37.7	58.8	39.0	4.7	4.0	48.4	70.0	63.7	9.0	54.2	33.3	37.4	61.6	57.6	30.1	31.7	32.4	52.8	49.0	27.8	40.2
Bilen <i>et al</i> [3]	66.4	59.3	42.7	20.4	21.3	63.4	74.3	59.6	21.1	58.2	14.0	38.5	49.5	60.0	19.8	39.2	41.7	30.1	50.2	44.1	43.7
Ren <i>et al</i> [25]	79.2	56.9	46.0	12.2	15.7	58.4	71.4	48.6	7.2	69.9	16.7	47.4	44.2	75.5	41.2	39.6	47.4	32.2	49.8	18.6	43.9
Wei <i>et al</i> [37]	67.3	63.3	61.3	22.7	8.5	64.8	57.0	80.5	9.4	49.0	22.5	72.6	73.8	69.0	7.2	15.0	35.3	54.7	75.0	29.4	46.9
Wang <i>et al</i> [33]	80.1	63.9	51.5	4.9	21.0	55.7	74.2	43.5	26.2	53.4	16.3	56.7	58.3	69.5	14.1	38.3	58.8	47.2	49.1	60.9	48.5
Cinbis <i>et al</i> [7]	67.1	66.1	49.8	34.5	23.3	68.9	83.5	44.1	27.7	71.8	49.0	48.0	65.2	79.3	37.4	42.9	65.2	51.9	62.8	46.2	54.2
Rochan <i>et al</i> [26]	78.5	63.3	66.3	56.3	19.6	82.2	74.7	69.1	22.4	72.3	31.0	62.9	74.9	78.3	48.6	29.3	64.5	36.2	75.8	69.5	58.8
pMWC (VGG-f)	59.7	67.1	60.3	46.4	51.2	68.8	75.9	57.9	40.4	77.3	21.5	64.6	65.2	74.7	67.3	41.6	77.1	48.0	60.9	60.9	59.3
pMWC (VGG-16)	64.7	58.4	60.3	54.1	52.0	71.0	79.2	63.8	43.1	71.6	40.5	64.4	72.1	84.9	69.5	45.3	75.0	51.1	66.3	71.1	62.9

Table 3. Co-localization CorLoc (%) of different methods on the PASCAL07 dataset.

Method	Aeroplane	Bird	Boat	Car	Cat	Cow	Dog	Horse	Motorbike	Train	Average
Prest <i>et al</i> [22]	51.7	17.5	34.4	34.7	22.3	17.9	13.5	26.7	41.2	25.0	28.5
Joulin <i>et al</i> [12]	25.1	31.2	27.8	38.5	41.2	28.4	33.9	35.6	23.1	25.0	30.9
Papazoglou <i>et al</i> [20]	65.4	67.3	38.9	65.2	46.3	40.2	65.3	48.4	39.0	25.0	50.1
Zhang <i>et al</i> [38]	75.8	60.8	43.7	71.1	46.5	54.6	55.5	54.9	42.4	35.8	54.1
Rochan <i>et al</i> [26]	56.0	30.1	39.6	85.7	24.7	87.8	55.6	60.2	61.8	51.7	55.3
pMWC (VGG-f)	48.5	74.4	52.8	61.6	59.4	69.3	71.4	68.5	73.6	43.0	62.3
pMWC (VGG-16)	44.3	68.6	56.7	63.5	50.0	70.7	71.2	75.9	73.8	55.5	63.0

Table 4. Co-localization CorLoc (%) of different methods on the YouTube-Objects dataset.

6.3 Experiments on the YouTube-Objects dataset

The Youtube-Objects dataset [13] is used for object co-localization in videos. The dataset contains videos collected from YouTube with 10 object classes. There are about 570,000 frames with 1,407 annotations in the first version of the dataset [22]. According to our knowledge, it is the largest available video dataset with bounding-box annotations on multiple classes. The individual video frames after decompression are used in our experiments to avoid possible confusion when applying different video decoders. We only perform object co-localization on video frames with ground-truth annotations, following the practice in [12]. No additional spatial-temporal information is utilized in our method. The co-localization accuracy of different methods on the YouTube-Objects dataset are summarized in Table 4. The experiments justify that the proposed object co-localization framework is also very effective for mining common objects in videos. Two other methods [38, 26] also utilize deep networks. The Youtube-Objects dataset comes with the test videos divided in 10 classes according to which dominant object is mostly present in them. Hence we construct 10 different graphs for this dataset.

7 Conclusion

In this paper, we present a novel framework to address the problem of object co-localization. It provides a practical and general solution for research and applications related to the MWC problem. Besides, deep learning based methods are

utilized to localize the candidates of the common objects and describe their visual characteristics. This makes it possible to better discriminate the inter-class similarities and identify the intra-class variations. Finally, a cycle elimination based restart strategy is proposed to guide the local search for the MWC. It successfully resolves the cycling issue in the optimization process. The experimental results on the object co-localization tasks demonstrate that our MWC solver is particularly suitable for graphs with high density. The proposed method shows significant improvements over several strong baselines.

Acknowledgements

Yi Fan is supported by the Natural Science Foundation of China (Nos.U1711263, U1811264, No.61463044, No. 61572234, No. 61672441) and the Shenzhen Basic Research Program (No. JCYJ20170818141325209). He is also donated a Titan Xp GPU card by NVIDIA corporation.

References

1. Adamczewski, K., Suh, Y., Mu Lee, K.: Discrete tabu search for graph matching. In: ICCV. pp. 109–117 (2015)
2. Alidaee, B., Glover, F., Kochenberger, G., Wang, H.: Solving the maximum edge weight clique problem via unconstrained quadratic programming. *European Journal of Operational Research* **181**(2), 592–597 (2007)
3. Bilen, H., Pedersoli, M., Tuytelaars, T.: Weakly supervised object detection with convex clustering. In: CVPR. pp. 1081–1089 (2015)
4. Chatfield, K., Simonyan, K., Vedaldi, A., Zisserman, A.: Return of the devil in the details: Delving deep into convolutional nets. In: BMVC (2014)
5. Chen, W., Chen, X., Zhang, J., Huang, K.: Beyond triplet loss: a deep quadruplet network for person re-identification. *arXiv* (2017)
6. Cho, M., Kwak, S., Schmid, C., Ponce, J.: Unsupervised object discovery and localization in the wild: Part-based matching with bottom-up region proposals. In: CVPR. pp. 1201–1210 (2015)
7. Cinbis, R.G., Verbeek, J., Schmid, C.: Weakly supervised object localization with multi-fold multiple instance learning. *TPAMI* (2016)
8. Everingham, M., Van Gool, L., Williams, C.K.I., Winn, J., Zisserman, A.: The PASCAL Visual Object Classes Challenge 2007 (VOC2007) Results (2007)
9. Everingham, M., Van Gool, L., Williams, C.K.I., Winn, J., Zisserman, A.: The PASCAL Visual Object Classes Challenge 2012 (VOC2012) Results (2012)
10. Fan, Y., Li, N., Li, C., Ma, Z., Latecki, L.J., Su, K.: Restart and random walk in local search for maximum vertex weight cliques. In: IJCAI (2017)
11. Gouveia, L., Martins, P.: Solving the maximum edge-weight clique problem in sparse graphs with compact formulations. *EURO Journal on Computational Optimization* **3**(1), 1–30 (2015)
12. Joulin, A., Tang, K., Fei-Fei, L.: Efficient image and video co-localization with frank-wolfe algorithm. In: ECCV. pp. 253–268 (2014)
13. Kalogeiton, V., Ferrari, V., Schmid, C.: Analysing domain shift factors between videos and images for object detection. *TPAMI* **38**(11), 2327–2334 (2016)

14. Kwak, S., Cho, M., Laptev, I., Ponce, J., Schmid, C.: Unsupervised object discovery and tracking in video collections. In: ICCV. pp. 3173–3181 (2015)
15. Li, Y., Liu, L., Shen, C., van den Hengel, A.: Image co-localization by mimicking a good detectors confidence score distribution. In: ECCV. pp. 19–34 (2016)
16. 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. pp. 740–755 (2014)
17. Ma, T., Latecki, L.J.: Maximum weight cliques with mutex constraints for video object segmentation. In: CVPR. pp. 670–677 (2012)
18. Mascia, F., Cilia, E., Brunato, M., Passerini, A.: Predicting structural and functional sites in proteins by searching for maximum-weight cliques. In: AAAI (2010)
19. McCreesh, C., Prosser, P., Simpson, K., Trimble, J.: On maximum weight clique algorithms, and how they are evaluated. In: ICPPCP. pp. 206–225 (2017)
20. Papazoglou, A., Ferrari, V.: Fast object segmentation in unconstrained video. In: ICCV. pp. 1777–1784 (2013)
21. Pattillo, J., Youssef, N., Butenko, S.: Clique relaxation models in social network analysis. In: Thai, M., Pardalos, P. (eds.) Handbook of Optimization in Complex Networks (2012)
22. Prest, A., Leistner, C., Civera, J., Schmid, C., Ferrari, V.: Learning object class detectors from weakly annotated video. In: CVPR. pp. 3282–3289 (2012)
23. Pullan, W.: Approximating the maximum vertex/edge weighted clique using local search. *Journal of Heuristics* **14**(2), 117–134 (2008)
24. Ren, S., He, K., Girshick, R., Sun, J.: Faster r-cnn: Towards real-time object detection with region proposal networks. *TPAMI* **39**(6), 1137–1149 (2017)
25. Ren, W., Huang, K., Tao, D., Tan, T.: Weakly supervised large scale object localization with multiple instance learning and bag splitting. *TPAMI* **38**(2), 405–416 (2016)
26. Rochan, M., Wang, Y.: Weakly supervised localization of novel objects using appearance transfer. In: CVPR. pp. 4315–4324 (2015)
27. Rosten, E., Drummond, T.: Machine learning for high-speed corner detection. In: ECCV. pp. 430–443 (2006)
28. Russakovsky, O., Deng, J., Su, H., Krause, J., Satheesh, S., Ma, S., Huang, Z., Karpathy, A., Khosla, A., Bernstein, M., et al.: Imagenet large scale visual recognition challenge. *IJCV* **115**(3), 211–252 (2015)
29. Simonyan, K., Zisserman, A.: Very deep convolutional networks for large-scale image recognition. *arXiv* (2014)
30. Tang, K., Joulin, A., Li, L.J., Fei-Fei, L.: Co-localization in real-world images. In: CVPR. pp. 1464–1471 (2014)
31. Uijlings, J.R., Van De Sande, K.E., Gevers, T., Smeulders, A.W.: Selective search for object recognition. *IJCV* **104**(2), 154–171 (2013)
32. Wang, C., Zhang, H., Yang, L., Cao, X., Xiong, H.: Multiple semantic matching on augmented n-partite graph for object co-segmentation. *TIP* **26**(12), 5825–5839 (2017)
33. Wang, C., Ren, W., Huang, K., Tan, T.: Weakly supervised object localization with latent category learning. In: ECCV. pp. 431–445 (2014)
34. Wang, L., Hua, G., Sukthankar, R., Xue, J., Niu, Z., Zheng, N.: Video object discovery and co-segmentation with extremely weak supervision. *TPAMI* **39**(10), 2074–2088 (2017)
35. Wang, X., Zhu, Z., Yao, C., Bai, X.: Relaxed multiple-instance svm with application to object discovery. In: ICCV. pp. 1224–1232 (2015)

36. Wang, Y., Cai, S., Yin, M.: Two efficient local search algorithms for maximum weight clique problem. In: AAAI. pp. 805–811 (2016)
37. Wei, X.S., Zhang, C.L., Li, Y., Xie, C.W., Wu, J., Shen, C., Zhou, Z.H.: Deep descriptor transforming for image co-localization. arXiv (2017)
38. Zhang, Y., Chen, X., Li, J., Wang, C., Xia, C.: Semantic object segmentation via detection in weakly labeled video. In: CVPR. pp. 3641–3649 (2015)
39. Zitnick, C.L., Dollár, P.: Edge boxes: Locating object proposals from edges. In: ECCV. pp. 391–405 (2014)