arXiv:1912.00501v1 [cs.CV] 1 Dec 2019

Interpreting Context of Images using Scene Graphs

Himangi Mittal[1], Ajith Abraham[2], and Anuja Arora[1]

[1] Department of Computer Science Engineering, Jaypee Institute of Information Technology, India
[2] Machine Intelligence Research Labs (MIR Labs), Auburn, WA 98071 USA, Washington United
himangimittal@gmail.com, ajith.abraham@ieee.org, anuja.arora@gmail.com

Abstract. Understanding a visual scene incorporates objects, relationships, and context. Traditional methods working on an image mostly focus on object detection and fail to capture the relationship between the objects. Relationships can give rich semantic information about the objects in a scene. The context can be conducive in comprehending an image since it will help us to perceive the relation between the objects and thus, give us a deeper insight into the image. Through this idea, our project delivers a model which focuses on finding the context present in an image by representing the image as a graph, where the nodes will the objects and edges will be the relation between them. The context is found using the visual and semantic cues which are further concatenated and given to the Support Vector Machines (SVM) to detect the relation between two objects. This presents us with the context of the image which can be further used in applications such as similar image retrieval, image captioning, or story generation.

Keywords: Scene Understanding · Context · Word2Vec · Convolution Neural Network

1 Introduction

Computer Vision has a number of applications which needs special attention of researchers such as semantic segmentation, object detection, classification, localization, and instance segmentation. The work attempted in the paper lies in the category of semantic segmentation. Semantic segmentation has two phases segmentation, detection of an object and semantic, is the prediction of context.

Understanding a visual scene is one of the primal goals of computer vision. Visual scene understanding includes numerous vision tasks at several semantic levels, including detecting and recognizing objects. In recent years, great progress has been made to build intelligent visual recognition systems. Object detection focuses on detecting all objects. Scene graph generation [1] [2] [3] [4] recognizes not only the objects but also their relationships. Such relationships can be represented by directed edges, which connect two objects as a combination of the

subject predicate - object. In contrast to the object detection methods, which just result in whether an object exists or not, a scene graph also helps in infusing context in the image. For example, there is a difference between a man feeding a horse and a man standing by a horse.

This rich semantic information has been largely unused by the recent models. In short, a scene graph is a visually grounded graph over the object instances in an image, where the edges depict their pairwise relationships. Once a scene graph is generated, it can be used for many applications. One such is to find an image based on the context by giving a query. Numerous methods for querying a model database are based on properties such as shape and keywords have been proposed, the majority of which are focused on searching for isolated objects. When a scene modeler searches for a new object, an implicit part of that search is a need to find objects that fit well within their scene. Using a scene graph to retrieve the images by finding context has a better performance than comparing the images on a pixel level. An extension to the above application is clustering of similar images. Recent methods cluster the image by calculating the pixel-to-pixel difference. This method does not generalize well and works on images which are highly similar. Also, this method may lead to speed and memory issues. The approach of scene graph infused with context can help to cluster the images even if there is a vast pixel difference. This method is also translation invariant, meaning, that a girl eating in the image can be anywhere in the image, but the context remains the same. Since this method uses semantic information, it enhances speed and memory.

The paper is structured in the following manner - Section II discusses the related work done in this direction, highlighting the scope of work to design a better solution. Importance and significance of work are discussed in section III. Section IV is about the dataset available and used to perform experiments. Section V discusses the solution approach followed by section VI which covers finding of object and context interpretation using scene graph. Finally, concluding remark and future scope is discussed in section VII.

2 Related Work

The complete work is can be divided into two tasks Object detection and Context interpretation. Hence, a plethora of papers have been studied to understand the various approaches defined by researchers in order to achieve an efficient and scalable outcome in both directions. Initially, in order to get an idea about deep learning models used in the field of computer vision, paper [6] is studied. This paper [6] covers the various deep learning models in the field of computer vision from about 210 research papers. It gives an overview of the deep learning models by dividing them into four categories - Convolutional Neural Networks, Restricted Boltzmann Machines, Autoencoder, and Sparse Coding. Additionally, their successes on a variety of computer vision tasks and challenges faced have also been discussed.

In 2016, Redmon et. al. [5] delivers a new approach, You Look Only Once (YOLO) for object detection by expressing it as a regression problem rather than a classification problem. It utilizes a single neural network which gives the bounding boxes coordinates and the confidence scores. Detection of context in images is an emerging application. Various methods ranging from scene graph to rich feature representation have been employed for the same. In 2018, Yang et. al. have developed a model Graph RCNN [4] which understands the context of the image by translating the image as a graph. A pipeline of object detection, pair pruning, GCN for context and SGGen+ has been employed. Similar sort of work is done by Fisher et. al. [7], they represent scenes that encode models and their semantic relationships. Then, they define a kernel between these relationship graphs to compare the common substructures of two graphs and capture the similarity between the scenes.

For effective semantic information extraction, Skipgram [8] model has been studies works for learning high quality distributed vector representation. In addition to this, several extensions [9] of Skipgram have been experimented with to improve the quality of vectors and training speed. Two models have been proposed in the work [10] which is an extension to Word2vec to improve the speed and time. Their architecture computes continuous vector representations of words from very large data sets. Large improvements have been observed in the accuracy at a much lower computational cost. The vectors are trained on a large dataset of Google for 1.6 billion words. In 2018, a new method deep structural ranking was introduced which described the interactions between objects to predict the relationship between a subject and an object. Liang et.al [11] makes use of rich representations of an image visual, spatial, and semantic representation. All of these representations are fused together and given to a model of structural ranking loss which predicts the positive and negative relationship between subject and object.

The work [12] aims to capture the interaction between different objects using a context-dependent diffusion network (CCDN). For the input to the model, two types of graphs are used - visual scene graph and semantic graph. The visual scene graph takes into account the visual information of object pair connections and the semantic graphs contain the rich information about the relationship between two objects. Once the features from visual and semantic graphs are taken, they are given as an input to a method called Ranking loss, which is a linear function. Yatskar et. al. work [13] is an extension to the predicate and relationship detection. It introduces a method where it focuses on the detection of a participant, the role of the participants and the activity of the participants. The model has coined the term "FrameNet" which works on a dataset containing 125,000 images, 11,000 objects, and 500 activities.

3 Importance and Significance of Work

This work is having its own importance and significance in varying application due to the following:

- An extension to the object detection by finding the underlying relationship between object and subject. Object detection merely works on the presence of the objects giving us partial information about the images. Context can give us the true meaning of the image.
- Classifies the image as similar on the basis of the underlying context. Object detection classifies the images as similar on the basis of the presence of specific objects. However, the images can be quite different than each other based on context. Incorporating the context will give a deeper insight into an image.
- If the context is employed on prepositions as well as verbs (future work), rich semantic information can be used to generate interesting captions and stories related to images.
- No pixel-to-pixel level similarity/clustering calculation. One of the applications of incorporating context is to find similar images. Conventional techniques involve pixel by pixel calculations, thus increasing the overhead. Scene graphs save time by considering the visual and semantic features.
- Useful in query processing, image retrieval, story generation, image captioning. Once the context is detected, it can be used in various applications like query processing in search engines, image retrieval using captioning [14][15], as well as story generation.

4 Datasets

Most famous datasets used for Scene understanding applications are MS-COCO [16], PASCAL VOC , and Visual Genome, and Visual Relationship Detection-VRD.

VRD dataset contains 5000 images, 100 object categories, and 70 predicates. It is most widely used for the relationship detection for an object pair in testing since it contains a decent amount of images. COCO [16] is large scale object detection, segmentation, and captioning dataset. This dataset is used in several applications- Object segmentation, Recognition in context, Super pixel stuff segmentation. It has over 330K images (200K labeled), and 80 object categories. Also, it has 5 captions per image which can be used for image captioning methods.

To perform the experiments, VRD dataset has been taken. Visual Relationship Detection (VRD) with Language Priors is a dataset developed by Stanford aiming to find the visual relationship and context in an image. The dataset contains 5000 images with 37,993 thousand relationships, 100 object categories and 70 predicate categories connecting those objects together. Originally, in the dataset, we are given a dictionary file of training and testing data which we convert into training set with 3030 images, test set of 955 images, and validation set of 750 images.

Statistics of the number of objects and visual relationships in every image is shown in Fig. 4 and Fig. 4 respectively. In Fig. 4, the file with an unusual number of 134 relationships in image '3683085307.jpg'.

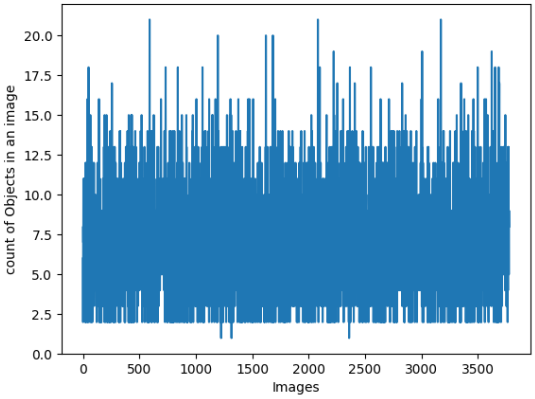

Fig.1. Statistics of the number of objects in images

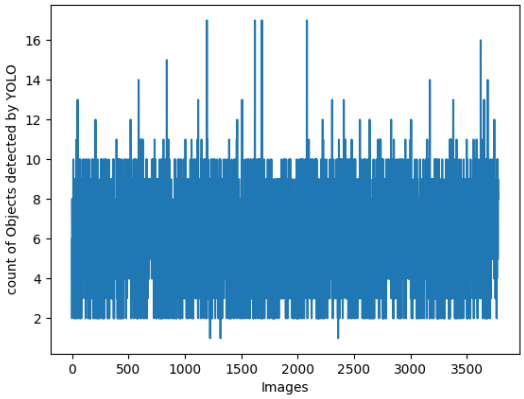

Fig.2. Statistics of the number of objects detected by YOLO

5 Solution Approach

The purpose of this project is to extend the object detection techniques to find context and hence, understand the underlying meaning of an image. Our idea uses 2D images of various scenes which have objects interacting with each other. The interaction can be in the form of prepositions like (above, on, beside, etc) or activity form like (cooking, eating, talking, etc). The model considers the scenery. The solution approach is basically divided into four modules. These modules are clearly depicted in Fig. 5. The first phase is the object detection for which YOLO object detector has been used. YOLO will provide an image with a bounding box of detected objects. This will be used to identify the semantic and visual features of the image. VGG-16 is used to generate visual features and Word2Vec is used for semantic feature identification. These features are concatenated and

given as input to a SVM which provides a probability distribution over all the predicate classes. (see Fig. 5 and Fig. 5.1).

Object Detection The first step of the solution approach is to detect the objects present in an input image. Recent research works have used various deep learning approaches and models. These are developed in order to achieve high efficiency and high accuracy for object detection. Approaches used in literature include YOLO [5], R-CNN [16], Fast-RCNN [17], Faster-RCNN [18], Mask-RCNN [20], and Single-Shot MultiBox Detector (SSD) [21].

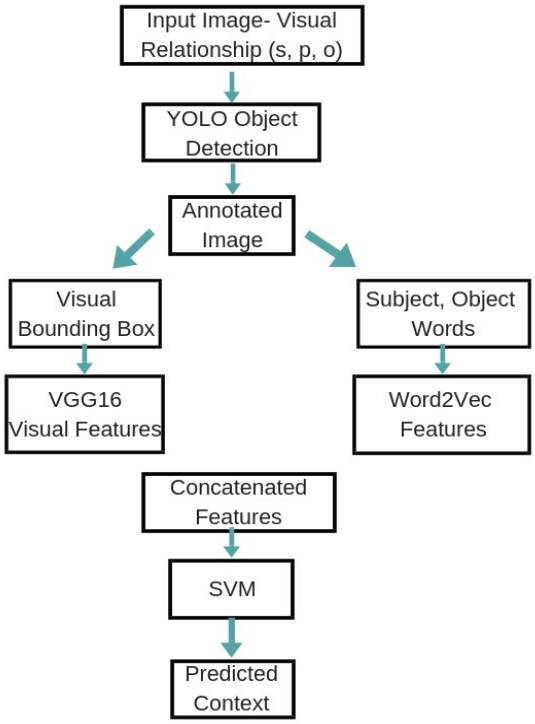

Fig.3. Flow Diagram of Solution Approach

Here, YOLO (You Only Look Once) has been used for object detection. It has an advantage that instead of using local filters, it looks at an image globally and delivers results. YOLO is very fast since it treats frame detection as a regression problem. The model consists of CNN similar to GoogleNet and instead of using residual block 1*1 convolutional layers are used. With 24 convolutional layers and pre-trained on ImageNet dataset, the model is trained for 135 epochs on PASCAL VOC dataset with a dropout of 0.5. Due to 1*1, the size of the prediction is the same as the feature space. The dimension of the feature space is in the format: 4 box coordinates, 1 objectness score, k class score for every box on the image. Object score represents the probability that an object is

contained inside a bounding box. Class confidences represent the probabilities of the detected object belonging to a particular class. YOLO uses softmax for class scores.

5.1 Semantic and Visual Embedding

Once the objects are detected, pairs for every object are created giving us nC2 number of visual relationships. For the visual features, the bounding box of subject and object are taken. The predicate is the intersection over union (IoU) of subject and object bounding boxes. All the three bounding boxes are concatenated and given to a VGG16 network with predicate word as the ground truth label. VGG16 is used for the classification task for the images. Its last layer provides good visual representations for the objects in an image. Hence, it is extracted to get visual relationship features for the concatenated bounding boxes. Further, for the semantic embedding, Word2Vec is used over the subject and object word. It is a powerful two layer neural network that can convert text into a vector representation. Word2Vec converts the subject and object word into a 300 sized feature representation which is concatenated and given to a neural network. The output layer before the application of activation function is extracted to get the semantic embedding for the visual relationship. The generated semantic embedding are stored in a dictionary format. The index is the object id and value is the embedding. We store the object, predicate and their embedding (found by word2vec) in the following formats shown in Table 1.

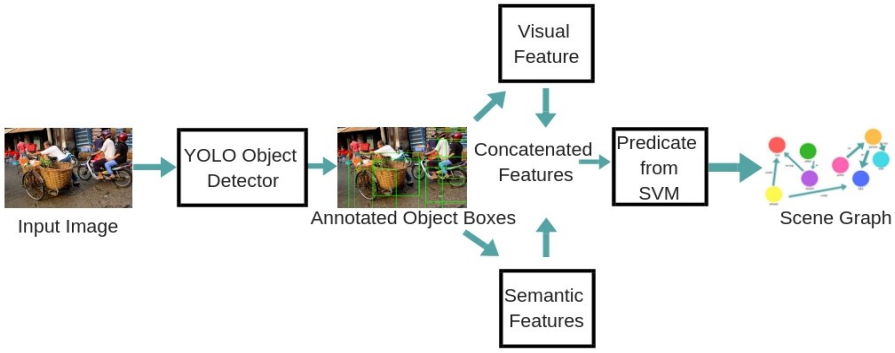

Fig.4. Phase Division for object detection and context interpretation

Considering that the visual and semantic embedding take the rich information about the image which is not limited to only object detection, but also to the semantic information present in an image, other information which can also be taken is the spatial feature representation which considers the location of an object in an image with respect to the other objects.

Table 1. File details of object, predicate, and embedding

File Name	Function	Dictionary Format
Objects_dict.pkl	Hashing of object	Index-Object Name
Predicate_dict.pkl	Hashing of predicates	Index-Predicate Name
Objects_embedding.pkl	Objects Word2Vec embedding	object name-word2vec embedding
Predicate_embedding.pkl	Predicates Word2Vec embedding	predicate name-word2vec embedding

5.2 Predicate Detection

The type of predicates in the dataset include the spatial predicates like above, beside, on, etc. Predicates can be of many types depicting the spatial context and activity context like cooking, eating, looking, etc. After the semantic and visual embeddings are extracted, the embedding is concatenated for a visual relationship in an image. Some other methods can also be used when using both the semantic and visual features which include, multiplying both the feature, however, this requires both the representation to be of the same size. The dataset includes around 70 predicates. Since the classes are quite distinct from each other, a decision boundary between the classes would serve as a good strategy to classify between the predicate classes and SVM is a powerful discriminative model to achieve this task. It is used as a classifier to give a class distribution probability over all the 70 classes. The class with the maximum probability is the predicted class. For the scene graph, top 3 predicates are taken. The predicate detected depicts the context shared between the subject and object and thus delivers the meaning of the image.

5.3 Scene Graph Generation

An image contains k number of visual relationships of the format (subject, predicate, object). The predicate was detected in the previous step. Now, the scene graph is generated with nodes as objects/subjects and edges as the predicate. Here, we use a directed scene graph so that there is a differentiation between subject and object.

For example: for a statement, a person eating food, the relationship format of (subject, predicate, object) would be (Person, eating, Food). Here, the person is the subject, food is the object, and eating is a predicate. If an undirected edge is used, this statement loses the distinctive property of the person being subject and food being object. The roles can be reversed due to undirected edges leading to erroneous relationships. Therefore, the use of directed edges is preferred. After the generation of the scene graph, it can be traversed accordingly to generate captions or summary of an image. The scene graph can also be termed as a context graph.

6 Findings

An outcome for a sample VRD dataset image is shown in Fig. 6. The image after YOLO is shown in Fig. 6 which shows the annotated image after YOLO object detection. The objects detected in the shown image in the boundary box are Person, wheel, cart, plant, bike, shirt, basket, and pants. Mean Average Precision (MAP) is taken as a performance measure to test the outcome. It considers the average precision for recall of the detected objects and is a popular metric for the object detectors. For the training set, YOLO had an object detection accuracy of 55 MAP.

Fig.5. Input image from the VRD-Dataset

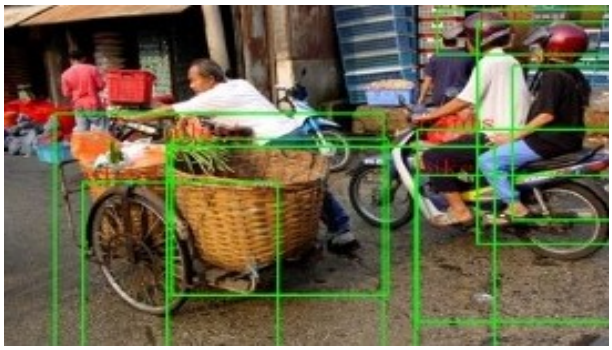

Fig.6. Annotated Image using YOLO Object Detector

Finally, a scene graph is generated based on these YOLO detected visual features and semantic features. The scene graph of Fig. 6 is depicted in Fig. 6. Relationships identified for which scene graph is formed are shown in Table 2 showing the scene description using the subject predicate-object relationship.

The loss in the Neural network for a semantic feature and CNN for the visual feature is shown in Fig. 6 and Fig. 6 respectively. It is clearly observable in Fig. 6 that the training loss dropped with every epoch. The validation, however, increased after the 50th epoch more than the training. The point where the validation loss increases the training loss depicts the point where the model starts to overfit. Hence, the weights of the network at the 50th epoch were taken

for further processing. One of the possible reason of overfitting can be attributed to the dataset being small. The model tries to fit to this small dataset and does not learn the ability to generalize well.

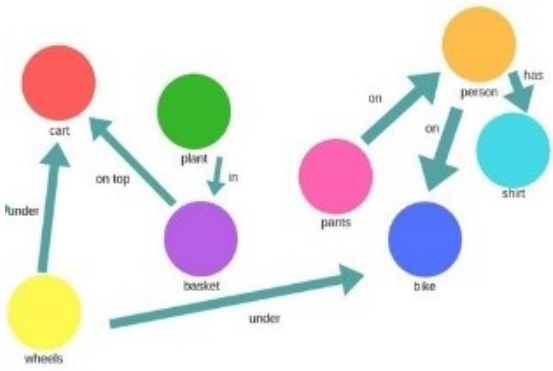

Fig.7. Scene Graph Result of the input image

Table 2. Details about scene, objects, and their relationships

Scene	Relationships (s, p, o)
Wheel under cart	(Wheel, under, cart)
Basket on top cart	(Basket, on top, cart)
Plant in basket	(Plant, in, basket)
Wheel under bike	(Wheel, under, bike)
Pants on person	(Pants, on, person)
Person on bike	(Person, on, bike)
Person has shirt	(Person, has, shirt)

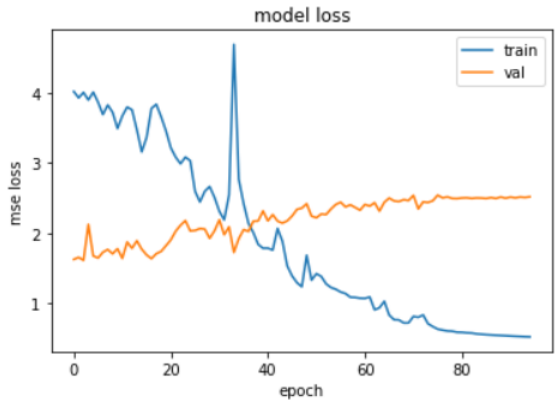

Fig.8. Training and Validation Loss for Visual Features

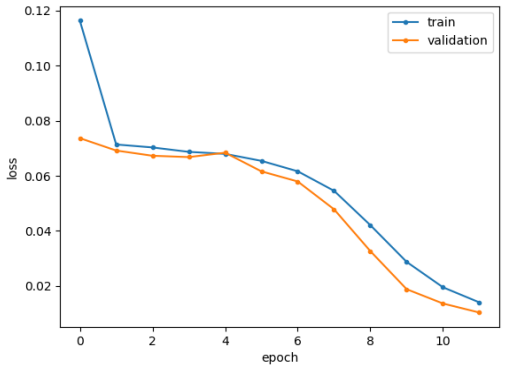

Fig.9. Training and Validation Loss for Semantic Features

The CNN and Neural Network were trained till they reached an accuracy of 95% and 99% respectively. The accuracy for predicate detection from SVM came out to be 60.57%. The SVM was run for a total for 100 epochs. In our previous approach of scene graph generation using Word2vec solely, the accuracy reached till 40% only. However, once we incorporated the visual features also, the accuracy increased to 60%

7 Conclusion and Future Scope

Our work leverages the techniques of object detection by finding out the context of the image in addition to the detected object. We are detecting the context from the visual and semantic features of an image. This is achieved by the application of deep learning models YOLO for object detection and Word2Vec for semantic feature representation generation. A neural network is used for the semantic feature of image and VGG16 for the visual feature generation. Context can be used to find out the subtle meaning of the image. Future work includes extending the context to verbs like cooking, eating, looking, etc since our work is covering only the preposition predicates such as on, above, etc. Moreover, in addition to the semantic and visual features, spatial features can be incorporated which will be helpful in determining the location of the objects. Lastly, better object detection models like Faster-RCNN can be employed for more accurate object detection in the first step because if an object is not detected in the first step, it can't be used for processing of visual relationships in the further stages.

References

1. Xu, Danfei, et al. "Scene graph generation by iterative message passing." Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. 2017.

2. Li, Yikang, et al. "Factorizable net: an efficient subgraph-based framework for scene graph generation." Proceedings of the European Conference on Computer Vision (ECCV). 2018.
3. Li, Yikang, et al. "Scene graph generation from objects, phrases and region captions." Proceedings of the IEEE International Conference on Computer Vision. 2017.
4. Yang, Jianwei, et al. "Graph r-cnn for scene graph generation." Proceedings of the European Conference on Computer Vision (ECCV). 2018.
5. Redmon, Joseph, et al. "You only look once: Unified, real-time object detection." Proceedings of the IEEE conference on computer vision and pattern recognition. 2016.
6. Guo, Yanming, et al. "Deep learning for visual understanding: A review." Neurocomputing 187 (2016): 27-48.
7. Fisher, Matthew, Manolis Savva, and Pat Hanrahan. "Characterizing structural relationships in scenes using graph kernels." ACM transactions on graphics (TOG). Vol. 30. No. 4. ACM, 2011.
8. Mikolov, Tomas, et al. "Efficient estimation of word representations in vector space." arXiv preprint arXiv:1301.3781 (2013).
9. Mikolov, Tomas, et al. "Distributed representations of words and phrases and their compositionality." Advances in neural information processing systems. 2013.
10. Mikolov, Tomas, et al. "Efficient estimation of word representations in vector space." arXiv preprint arXiv:1301.3781 (2013).
11. Liang, Kongming, et al. "Visual relationship detection with deep structural ranking." Thirty-Second AAAI Conference on Artificial Intelligence. 2018.
12. Cui, Zhen, et al. "Context-dependent diffusion network for visual relationship detection." 2018 ACM Multimedia Conference on Multimedia Conference. ACM, 2018.
13. Yatskar, Mark, Luke Zettlemoyer, and Ali Farhadi. "Situation recognition: Visual semantic role labeling for image understanding." Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. 2016.
14. Chen, Xinlei, and C. Lawrence Zitnick. "Mind's eye: A recurrent visual representation for image caption generation." Proceedings of the IEEE conference on computer vision and pattern recognition. 2015.
15. Gao, Lizhao, Bo Wang, and Wenmin Wang. "Image Captioning with Scene-graph Based Semantic Concepts." Proceedings of the 2018 10th International Conference on Machine Learning and Computing. ACM, 2018.
16. Lin, Tsung-Yi, et al. "Microsoft coco: Common objects in context." European conference on computer vision. Springer, Cham, 2014.
17. Girshick, Ross, et al. "Rich feature hierarchies for accurate object detection and semantic segmentation." Proceedings of the IEEE conference on computer vision and pattern recognition. 2014.
18. Girshick, Ross. "Fast r-cnn." Proceedings of the IEEE international conference on computer vision. 2015.
19. Ren, Shaoqing, et al. "Faster r-cnn: Towards real-time object detection with region proposal networks." Advances in neural information processing systems. 2015.
20. He, Kaiming, et al. "Mask r-cnn." Proceedings of the IEEE international conference on computer vision. 2017.
21. Liu, Wei, et al. "Ssd: Single shot multibox detector." European conference on computer vision. Springer, Cham, 2016.