

Iterative Shrinking for Referring Expression Grounding Using Deep Reinforcement Learning

Mingjie Sun^{1,2}, Jimin Xiao^{1,*}, Eng Gee Lim¹

¹Xi'an Jiaotong-Liverpool University, ²University of Liverpool

mingjie.sun@liverpool.ac.uk, {jimin.xiao,enggee.lim}@xjtlu.edu.cn

Abstract

In this paper, we are tackling the proposal-free referring expression grounding task, aiming at localizing the target object according to a query sentence, without relying on off-the-shelf object proposals. Existing proposal-free methods employ a query-image matching branch to select the highest-score point in the image feature map as the target box center, with its width and height predicted by another branch. Such methods, however, fail to utilize the contextual relation between the target and reference objects, and lack interpretability on its reasoning procedure. To solve these problems, we propose an iterative shrinking mechanism to localize the target, where the shrinking direction is decided by a reinforcement learning agent, with all contents within the current image patch comprehensively considered. Besides, the sequential shrinking processes enable to demonstrate the reasoning about how to iteratively find the target. Experiments show that the proposed method boosts the accuracy by 4.32% against the previous state-of-the-art (SOTA) method on the RefCOCOg dataset, where query sentences are long and complex with many targets referred by other reference objects.

1. Introduction

The aim of referring expression grounding (REG) is to recognize and localize the target object in an image according to its query sentence (referring expression), which requires joint comprehension in both visual and linguistic domains. REG is a fundamental multi-modality task, serving as the basis for many downstream tasks, including the visual question answering [15, 4, 17, 54, 20, 1, 48], image caption [9, 5, 7, 34, 65, 49, 42], image-text matching [25, 65, 3, 64, 6, 51, 21], etc. It attracts much attention from

*corresponding author

¹The work was supported by National Natural Science Foundation of China under 61972323 and Key Program Special Fund in XJTLU under KSF-T-02, KSF-P-02.

<https://github.com/insomnia94/ISREG>

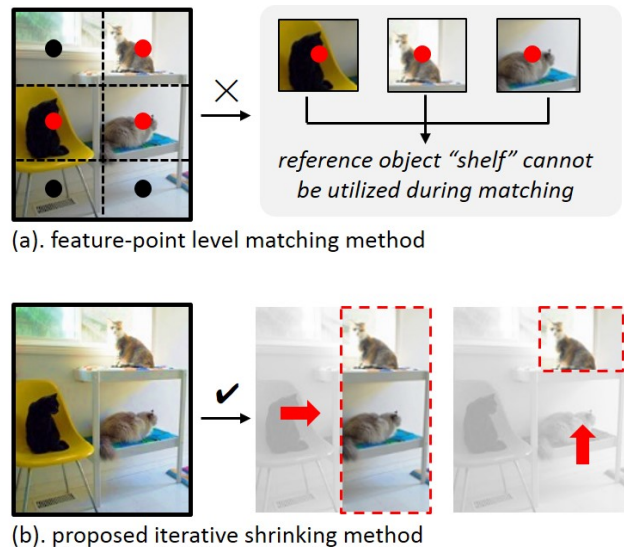


Figure 1. Illustration of the difference between (a) existing feature-point level matching methods, and (b) the proposed iterative shrinking method. The query sentence for the target in this image is “the cat above the shelf”. In (a), dots depict the feature points within the image feature map, with the surrounding area representing its corresponding image region. Red dots indicate the feature points whose corresponding image regions may cover the target (i.e., the cat on the shelf) or distracting objects (i.e., cats on other places), while black dots are unrelated feature points. It can be observed that there is no such feature point whose image region can cover both the target object (i.e., cat) and the reference object (i.e., shelf). In (b), for each iteration, the area in the dashed box indicates the image patch after shrinking, and the red arrow depicts the shrinking direction, which is decided by comprehensively considering all contents in the image patch before shrinking.

both computer vision (CV) and natural language processing (NLP) areas in recent years.

Conventional methods formulate REG as a region-retrieval problem [61, 63, 56, 53, 26, 27, 10], with proposals of all candidate objects provided in advance. The candidate proposals can be either provided from the bounding box

ground-truth or generated by a pre-trained object detector (e.g., Faster RCNN [38]). The matching networks of these REG models predict the similarity score between the query and each proposal, with the highest-score proposal selected as the target object. In this way, conventional REG methods highly rely on the ground-truth bounding boxes or an accurate object detector which requires a lot of extra data to train it in advance. To overcome these shortcomings, some proposal-free methods attempt to predict the REG result without object proposals.

All existing proposal-free REG methods [22, 59, 41, 58] directly follow the pipeline of one-stage detector (e.g., YOLOv3 [37]) and adopt a two-branch style network, with the first branch to calculate the similarity score between the query sentence and each feature point within the image feature map, and the second branch to generate its bounding box coordinates. The predicted bounding box of the feature point with the highest score is viewed as the final result.

These methods, with the matching process conducted between the query sentence and each image feature point, perform well for simple queries with the target described by its own attribute (e.g., “man in blue”), but it is difficult for them to deal with complex queries, especially when the target is referred by another reference object. As can be observed from Fig.1.(a), since there is no such a feature point whose corresponding image region covers both the target “cat” and the reference object “shelf”, the contextual relation between the target and the reference object, serving as the key to distinguish the target from other distracting objects, cannot be fully utilized in such a matching process. Thus, the REG performance is greatly compromised.

Another shortcoming of these methods is lack of interpretability. As the matching network predicts the final matching scores of all feature points in one step, its inner reasoning procedure is hidden and invisible. Thus, if the matching model fails, it is formidable to analyze the causes.

To tackle these problems simultaneously, we formalize REG as a sequence of image-level shrinking processes. Within each iteration, the image shrinks along a certain direction, with a non-target image region removed. After several shrinking iterations, only the target image region remains, and it is viewed as the final result. The shrinking direction in each iteration is predicted by a trainable network. As the optimal shrinking direction is uncertain for each iteration, conventional supervised training is not suited. We propose to model it as a Markov decision process [14] and adopt reinforcement learning (RL) to tackle it. RL only requires a “feedback” or “reward” after each shrinking step, rather than the exact supervision label. Besides, RL considers not only the current reward but also the potential reward in future, which further improves the performance.

Our proposed method can better utilize the contextual relation between the target and reference objects because the

shrinking direction is decided by comprehensively considering all objects within the current image patch. For an instance, as shown in Fig.1.(b), in the first iteration, the image region where a black cat lies on the chair is removed (i.e., shrinking towards right), because “chair” does not occur in the query. In the second iteration, after analyzing the spatial relation between each candidate cat and the reference object “shelf”, it is decided to remove the image region where a cat lies inside the shelf (i.e., shrinking towards top), due to its mismatching spatial relation against the query. Ultimately, only the image region accurately covering the target cat remains. As the relation information in a query is fully utilized, the proposed method can better deal with complex queries and images (e.g., RefCOCOg [30]), whilst existing matching methods fail to handle such cases (Fig.1.(a)).

To sum up, this paper has three main contributions:

- We make the earliest attempt to solve REG as a conditional decision-making process and build the first RL-based REG framework. Thanks to these iterative decisions, the reasoning about how to localize the target can be visualized step by step.
- We formalize REG as a sequence of image-level shrinking processes, with the shrinking direction in each iteration decided by an RL agent after comprehensively considering all objects in the image patch, allowing the relation information in a query to be fully utilized.
- The proposed method boosts the accuracy by 4.32% against the previous SOTA method [22] on the RefCOCOg dataset, where query sentences are long and complex, with many targets referred by other reference objects. An average accuracy gain of 1.33% on RefCOCOg, RefCOCO+ and RefCOCO is also achieved.

2. Related Work

2.1. Referring Expression Grounding

REG can be classified into two categories, including proposal-based and proposal-free. For many early proposal-based REG methods, the entire query sentence is directly encoded through a single language embedding network (e.g., LSTM [12]), without considering the variance among different types of query sentences [40, 52, 30, 62]. To solve this problem, Yu *et al.* [61] propose an attention mechanism to decompose the query sentence into three linguistic components, related to subject appearance, location and relationship to other objects. The ultimate matching result is generated after comprehensively considering all three components. To prevent the attention mechanism from only focusing on the most dominant features of both modalities,

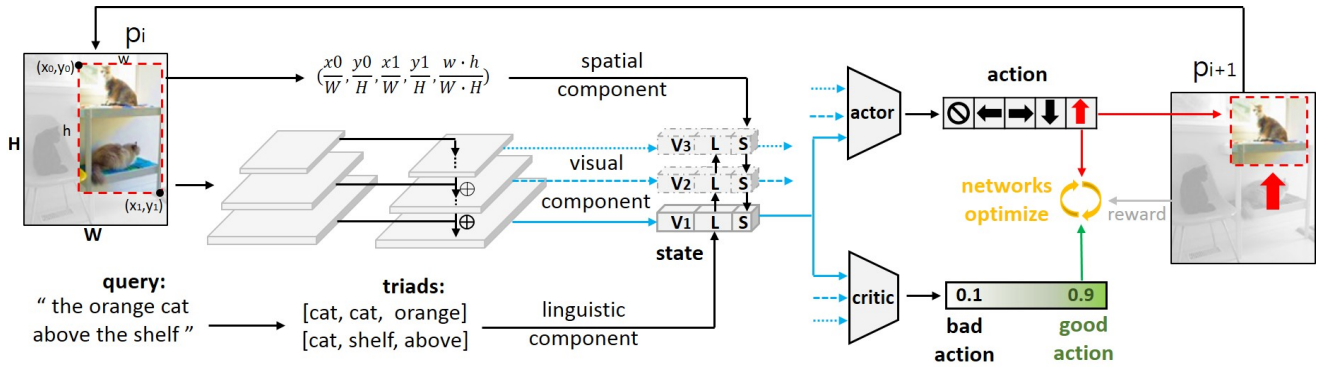


Figure 2. Architecture of our method where the blue lines indicate the states corresponding to different scale levels, and the yellow circle represents the optimizing of the actor and critic networks within each iteration.

an erasing mechanism is proposed in [28], where the most dominant linguistic or visual information is discarded to discover more complementary linguistic-visual correspondences.

In the proposal-free REG setting, Sadhu *et al.* [59] combine the linguistic feature of the query sentence and the visual feature from the image to generate a multi-modal feature map, and predict the binary classification scores (foreground or background) for all points and their corresponding bounding box coordinates. RCCF [22] adopts a correlation filter to calculate the similarity between the query sentence and each point in the image feature map, and the point with the highest similarity score is viewed as the center of the target object. Then, the size and offset are predicted by another network to generate the final bounding box for the target.

The aforementioned proposal-free REG methods [59, 22] only consider a local image region and ignore the contextual object relations. Differently, we propose an iterative shrinking mechanism to predict the target object location, where the shrinking direction within each iteration is decided by comprehensively considering all objects in the current image, allowing the relation information in a query to be better utilized.

2.2. Deep Reinforcement Learning

Deep reinforcement learning algorithms have achieved great success on sequential decision-making problems. “Actor-critic” [19] is a popular RL framework that comes from several previous RL frameworks including deep Q-learning [33] and policy gradient [47] which are value-based and policy-based strategies, respectively.

RL has been applied to many CV tasks [45, 46], especially multi-modal problems. Rennie *et al.* [39] apply the policy gradient to train the image caption network based on its baseline method [55]. Specifically, the image caption network is viewed as an RL agent, whose action is to predict the next word. The state of the RL agent is the cells and hid-

den variables of the LSTM, as well as their corresponding attention weights. The reward function is designed according to the sentence-level evaluation metric, CIDEr score [50], representing the difference between the generated sentence and the ground-truth one. Yu *et al.* [63] adopt RL to guide the referring sentence generator towards less ambiguous expression. Different from [39], the reward function is designed according to the similarity score between the generated sentence and the object image patch, calculated by a pre-trained region-language matching network.

However, to the best of our knowledge, RL has never been used to tackle the REG task. All existing REG methods follow the static region-sentence matching strategy, leaving no space to make any sequential decisions. In this paper, we abandon the traditional matching strategy, and make the first attempt to formulate REG as a sequential shrinking processes, where an RL agent is trained to predict the shrinking direction within each iteration.

3. Methodology

3.1. Overview

Given the image img and the query sentence q , the target object is localized through a sequence of shrinking iterations, with i indicating the index of each iteration. Specifically, for a certain iteration i , the current image patch p_i shrinks along a certain direction, decided by the RL agent, into a smaller patch p_{i+1} , with a non-target region eliminated. After several sequential iterations, all non-target regions are removed and only the region exactly covering the target object remains, which is viewed as the final REG result. Note that $p_1 = img$ at the beginning.

In terms of the shrinking direction selection, we adopt an RL agent to make the decision within each iteration, and view the aforementioned shrinking procedure as a Markov decision process (MDP) [14]. Specifically, within a certain iteration i , the state s_i is generated according to the information embedded in p_i and q (Section 3.2). Then, according

to s_i , the RL agent predicts an action a_i (shrinking direction or stop). After a_i is executed, a reward r_i is calculated as a “feedback” of a_i (Section 3.3). Ultimately, the RL agent is trained through an actor-critic RL framework, where the RL agent consists of two components, including an actor network to predict the action, and a critic network to evaluate the predicted action (Section 3.4). The architecture of our proposed method is shown in Fig.2.

3.2. Agent State

In conventional MDP, states indicate the observations that an agent receives from the external environment. Specifically to our method, for a certain image patch p_i within iteration i , the state s_i is the input passed into the RL agent to predict the action, consisting of three components, including the linguistic feature f_q^l extracted from q , visual feature $f_{p_i}^v$ extracted from p_i , and spatial feature $f_{p_i}^s$ extracted from the relative position of p_i within img .

3.2.1 Linguistic Component

To extract the linguistic feature f_q^l , q is first parsed into multiple discriminative triads $\{t_k\}_{k=1}^M$ [44]. Each triad represents a piece of discriminative information to distinguish the target from the distracting or reference objects. As can be observed from Table 1, a discriminative triad t_k consists of three components, including a target unit t_k^t , a reference unit t_k^r and a discriminative unit t_k^d , denoting the target object, the reference object against the target object and the discriminative relation between them, respectively. Note that, the discriminative unit t_k^d is defined in a broad sense to take full advantage of the potential discriminative information hidden in queries. As for an instance, when it comes to “orange cat” in Table 1, this query implies that there are cats in other colors with a high probability. Thus, “cat” is set as its reference unit to represent the distracting cats in other colors.

For its implementation, an off-the-shelf NLP processing toolbox (e.g., Stanford CoreNLP [2] or Spacy [13]) is employed to analyze the tree structure of each query, as well as the dependency label of each word in the query. In general, triads are generated within two steps. The first step is to extract the rightmost Normal Noun (NN) of the bottom-left Noun Phrase (NP) in the sentence tree as the target unit shared by all triads. The second step is to generate the reference and discriminative units using different parsing patterns. As for an instance, the subject-relation-object phrase, such as “the man holding a cat”, firstly the second component of a particular dependency, whose first component is the same as the target unit, and the dependency type is nsubj, is viewed as a discriminative unit. Then, the second component of another dependency, whose first component is the same as that of the current discriminative unit, is

Table 1. Common queries and their corresponding discriminative triads on the RefCOCOg, RefCOCO+ and RefCOCO datasets. **ID** is the index of a discriminative triad for a query, as a complex query may generate multiple discriminative triads. **T/R/D** indicates the target, reference, and discriminative unit, respectively.

Query	ID	T	R	D
“lady”	1	lady	lady	SELF
“left”	1	UKN	UKN	left
“left lady”	1	lady	lady	left
“orange cat”	1	cat	cat	orange
“cat above a shelf”	1	cat	shelf	above
“lady holding a cat”	1	lady	cat	holding
“the left lady in white holding an orange cat and standing on a table”	1	lady	lady	left
	2	lady	lady	white
	3	lady	table	on
	4	lady	cat	holding
	5	cat	cat	orange

viewed as its corresponding reference unit.

Then, a pre-trained word vector generation method (e.g., Word2vec [31] or Glove [35]) is adopted to extract the linguistic embedding $f_{t_k^t}^l, f_{t_k^r}^l, f_{t_k^d}^l$, corresponding to t_k^t, t_k^r, t_k^d , respectively. The linguistic feature $f_{t_k}^l$ of a certain triad t_k is the concatenation of its three components’ linguistic features:

$$f_{t_k}^l = f_{t_k^t}^l \oplus f_{t_k^r}^l \oplus f_{t_k^d}^l, \quad (1)$$

where \oplus indicates the concatenation operation. Similarly, another concatenation operation is further conducted on all triads’ linguistic features:

$$f_q^l = f_{t_1}^l \oplus f_{t_2}^l \oplus \dots \oplus f_{t_M}^l, \quad (2)$$

and f_q^l represents the ultimate linguistic feature of q . Note that, if the total number of the triads parsed from a query is less than M , the first triad feature $f_{t_1}^l$ is repeated until the triad number reaches M .

3.2.2 Visual and Spatial Components

The visual feature of the image patch p_i is extracted through an off-the-shell convolutional network (e.g., VGG [43] or ResNet [11]) and the feature pyramid network (FPN) [23]. Specifically, the convolutional network, e.g., ResNet [11], first generates three raw visual feature maps, $f_{p_i}^{v_1}, f_{p_i}^{v_2}, f_{p_i}^{v_3}$, with different scales of $[28 \times 28 \times 512], [14 \times 14 \times 1024]$ and $[7 \times 7 \times 2048]$, respectively. Then, the up-sampling, fusing and average-pooling operations of FPN [23] are conducted on these three raw feature maps, which generates three one-hot vectors with the same dimension (D_v), $f_{p_i}^{v_1}, f_{p_i}^{v_2}$ and $f_{p_i}^{v_3}$, as the corresponding visual feature for each scale level.




Figure 3. Illustration of our action design. $\{a^1, a^2, a^3, a^4\}$ corresponds to image shrinking toward 4 different directions, respectively. a^5 corresponds to the STOP action.

The spatial feature $f_{p_i}^s$ of the image patch p_i is encoded by a 5-D vector, $[\frac{x_{tl}}{W}, \frac{y_{tl}}{H}, \frac{x_{br}}{W}, \frac{y_{br}}{H}, \frac{w \cdot h}{W \cdot H}]$, where $[x_{tl}, y_{tl}]$ and $[x_{br}, y_{br}]$ refer to the top-left and bottom-right coordinates of p_i within img , respectively. w and h indicate the width and height of p_i , while W and H denote the width and height of img . $f_{p_i}^s$ can provide the relative position of p_i within img , through which the RL agent can infer the relative location of the target object within img .

The ultimate state s_i consists of three sub-states s_i^c with $c \in \{1, 2, 3\}$, each of which corresponds to a certain scale level c . Each sub-state is designed as the concatenation of the corresponding linguistic, visual and spatial features:

$$s_i^c = f_q^l \oplus f_{p_i}^{v_c} \oplus f_{p_i}^s. \quad (3)$$

3.3. Action and Reward

In MDP, after observing the environment, an agent takes actions to interact with the external environment. A reward is given to provide the feedback of the conducted action. As shown in Fig.3, within a certain iteration i , the RL agent selects an action a_i among the action pool A , consisting of 5 candidate actions, $\{a^1, a^2, a^3, a^4, a^5\}$. The first 4 candidate actions, $\{a^1, a^2, a^3, a^4\}$, indicate the top side of p_i shrinks towards down, the bottom side of p_i shrinks towards up, the left side of p_i shrinks towards right and the right side of p_i shrinks towards left, respectively. a_5 indicates the end of the shrinking process (i.e., $p_i = p_{i+1}$), and the current image patch is viewed as the final REG result.

To naturally introduce the shrinking stride decay in an episode, the value of the shrinking stride is decided according to the width and height of the image patch p_i before shrinking:

$$\begin{aligned} \hat{y}_{tl} &= y_{tl} + \alpha(y_{br} - y_{tl}), \text{ if } a_i = a^1 \\ \hat{y}_{br} &= y_{br} - \alpha(y_{br} - y_{tl}), \text{ if } a_i = a^2 \\ \hat{x}_{tl} &= x_{tl} + \alpha(x_{br} - x_{tl}), \text{ if } a_i = a^3 \\ \hat{x}_{br} &= x_{br} - \alpha(x_{br} - x_{tl}), \text{ if } a_i = a^4, \end{aligned} \quad (4)$$

where $[x_{tl}, y_{tl}]$ and $[x_{br}, y_{br}]$ indicate the coordinates of the top-left and bottom-right points of p_i within img . $[\hat{x}_{tl}, \hat{y}_{tl}]$ and $[\hat{x}_{br}, \hat{y}_{br}]$ indicate the coordinates of the top-left and bottom-right points of p_{i+1} within in img . In this way,

at the early stage of an episode, relatively large shrinking strides are adopted to quickly “move” towards to the target. While in the late stage, relatively small shrinking strides are adopted to meticulously localize the accurate position of the target.

To predict such an action a_i , the state s_i^c for each scale level is passed into the RL agent respectively, generating three action probability vectors, with the average one serving as the final probability vector. The action with the highest probability is selected as a_i .

The reward r_i is designed according to the overlap between the image patch after shrinking and the ground-truth bounding box, as follows:

$$r_i = \begin{cases} 0 & IoU < 0.3 \text{ or } \Delta_{IoU} \leq 0 \\ 1 & 0.3 \leq IoU < 0.5 \text{ and } 0 < \Delta_{IoU} \\ 10 & 0.5 \leq IoU \text{ and } 0 < \Delta_{IoU} \end{cases}, \quad (5)$$

where IoU refers to the intersection-over-union between the image region after shrinking and the ground-truth bounding box of the target, and Δ_{IoU} indicates the IoU difference between the current iteration and the previous iteration. We set major difference between the good and bad action’s rewards, so as to facilitate the training of the RL agent.

3.4. Actor-critic Training

The proposed RL agent is designed and optimized under the “actor-critic” framework [19], consisting of two networks, including the actor network to predict the action among the candidate action pool, and the critic network to predict the score of a certain state, which is further utilized to evaluate the quality of the predicted action. During each iteration of training, the actor and critic networks are optimized in order, while in the inference stage, only the actor network is used to predict the action to localize the target.

Specifically, the actor network is first optimized in a policy-based way:

$$\theta = \theta' + l_a \nabla(\log \pi_{\theta'}(s_i, a_i)) A(s_i, a_i), \quad (6)$$

where

$$A(s_i, a_i) = \delta_i = r_i + \gamma V_{w'}(s_{i+1}) - V_{w'}(s_i). \quad (7)$$

In Eq.(6) and Eq.(7), θ' and θ indicate the weights of the actor network before and after the update. l_a is the learning rate corresponding to the actor network. $\pi(s, a)$ is the actor network’s policy function to predict the probability of taking action a for state s . $A(s, a)$ is the advantage function, which is equal to the TD error function in the “actor-critic” framework [32]. r_i is the reward calculated using Eq.(5), and γ refers to the discount factor. w' is the weight of the critic network before the update of the critic network. $V(s)$

Table 2. Accuracy comparison against other REG methods. “PF” indicates proposal-free REG methods. Specifically, methods on the top half of the table are not proposal-free, while methods on the bottom half are proposal-free. “Average” is the average accuracy of RefCOCOg, RefCOCO+, and RefCOCO, where mean accuracy of testA and testB is used for RefCOCO+ and RefCOCO. The method with the highest accuracy is in bold.

Method	PF	Visual Extractor	Average	RefCOCOg	RefCOCO+		RefCOCO	
				test	testA	testB	testA	testB
SLR[63] _{CVPR17}		ResNet-101	61.25	59.63	60.74	48.80	73.71	64.96
DGA[57] _{ICCV19}		VGG-16	65.26	63.28	69.07	51.99	78.42	65.53
MAttNet[61] _{CVPR18}		ResNet-101	68.33	67.01	70.26	56.00	80.43	69.28
FAOAVG[59] _{ICCV19}		DarkNet-53	61.79	57.45	60.56	52.86	74.81	67.59
MCN[29] _{CVPR20}	✓	VGG-16	65.69	62.29	65.24	54.26	76.97	73.09
RCCF[22] _{CVPR20}		DLA-34	67.30	65.73	70.35	56.32	78.93	66.73
Ours		ResNet-101	68.63	70.05	71.05	58.25	74.27	68.10

* RefCOCO results of RCCF are from Table 4, setting 7 in [22], where the feature backbone is obtained without the detection pretraining using the additional classification labels of the 80 categories in MS-COCO dataset. As the results of RefCOCOg and RefCOCO+ are not provided in Table 4, results from Table 3 in [22] are used instead, but the comparison is actually unfavorable for the proposed method. The darknet-based MCN [29] is not included in this table for the same reason.

indicates the predicted score for state s , generated by the critic network.

After the actor network is updated, the critic network is optimized in a value-based way:

$$w = w' + l_c \delta_i \nabla_{w'} V_{w'}(s_i). \quad (8)$$

In Eq.(8), w' and w indicate the weights of the critic network before and after the optimization. l_c is the learning rate corresponding to the critic network. δ_i is the TD error function calculated in Eq.(7).

In this way, the joint optimization of the actor and critic networks can avoid the disadvantages of both value-based and policy-based methods during training. In other words, the proposed RL agent can be optimized after each iteration, rather than waiting until the end of each episode, which dramatically speeds up the training process yet maintains training stability.

3.5. Implementation Details

To generate the state, Stanford CoreNLP [2] is adopted to parse the query sentence into triads, and Glove word vector [35] is adopted to generate the linguistic feature of each triad. The triad number M for each query is set as 2, because a too small M cannot contain enough information from queries, while an extremely large M may greatly increase the network complexity. Resnet101 [11] is used to generate the raw visual feature of the image patch, and FPN [23] is adopted to generate the fused visual feature, with D_v set as 512. In terms of the action, α is set as 0.2 to generate a moderate shrinking stride. During the training, γ is set as 0.9, and the learning rate is set as 1e-7 for both the

Table 3. Performance comparison of our method against other REG methods on the RefCLEF dataset. The method with the highest accuracy is in bold.

Method	Accuracy
CITE[36] _{ECCV18}	34.13
IGOP[60] _{NIPS17}	34.70
RCCF[22] _{CVPR20}	63.79
Ours	65.48

actor and critic networks. ADAM [18] is adopted as the optimizer. Training is conducted on one Titan RTX GPU for around 800,000 episodes.

Similar as [22], to obtain a more precise bounding box for a target, we employ a post-processing box refinement model after the shrinking process, where the input is img and the predicted bounding box from the shrinking process, and the output is the new refined bounding box. Such box refinement network is trained using the L2 loss between the coordinates of a ground-truth box in MS-COCO [24] and the coordinates of corresponding sampled box obtained by adding Gaussian noise to it.

4. Experiments

4.1. Datasets and Metrics

The proposed method is evaluated on four common datasets including RefCOCO [62], RefCOCO+ [62], RefCOCOg [30] and RefCLEF [16]. The first three datasets are collected from MS-COCO [24]. The main difference between RefCOCO+ and RefCOCO is that the former one

Table 4. Ablation studies on RefCOCO (testA) and RefCOCOg (test) datasets. The method with the highest accuracy is in bold.

Setting	RefCOCO	RefCOCOg
w/o triad	67.52	60.17
w/o multi-scale	73.33	66.35
w/o spatial feature	62.42	54.88
fixed stride	45.70	39.46
w/o refinement	70.18	65.62
supervised training	63.84	61.36
ours	74.27	70.05

forbids the absolute location words when generating the query sentences, requiring more attention to the appearance discrimination. RefCOCOg contains longer expressions for both appearance and locations. Referring sentences in RefCLEF are annotated casually, so it is normally used as an auxiliary dataset. The metric used to evaluate the accuracy of the proposed method is similar to the object detection task, where the intersection-over-union (IoU) between the predicted bounding box and the ground-truth one is calculated. Then, the one with IoU higher than 0.5 is treated as correct.

4.2. Comparison with State-of-the-Art

To evaluate our proposed method, we compare it with other SOTA REG methods, which are divided into two groups. The first group, shown in the top half of Table 2, consists of REG methods with object proposals, including SLR [63], DGA [57] and MAttNet [61]. The second group, shown in the bottom half of Table 2, is composed of proposal-free REG methods, including FAOAVG [59], MCN [29] and RCCF [22]. As can be observed from Table 2, the proposed method achieves a large accuracy gain (4.32%) over the previous SOTA method, RCCF [22], on the test set of the RefCOCOg dataset. When the backbone is replaced by VGG, the proposed method achieves 66.32% on this dataset. The cause of such big accuracy gain is that, queries in RefCOCOg are much longer and more complex than queries in other datasets, with many targets referred by other neighboring reference objects, namely “targets with reference”. Thus, the key to localize such targets is to take full advantage of the relation information between the target and reference objects in queries, which is difficult for the feature-point level matching mechanism [59, 29, 22]. Nevertheless, the image-level shrinking mechanism allows the proposed method to consider all objects within the image patch, and better utilize such contextual relation information, so as to achieve higher accuracy.

Then, similar accuracy improvement can be seen on RefCOCO+ (1.32%) and RefClef (1.69%) datasets, where many queries are complicated as well, as shown in Table 2

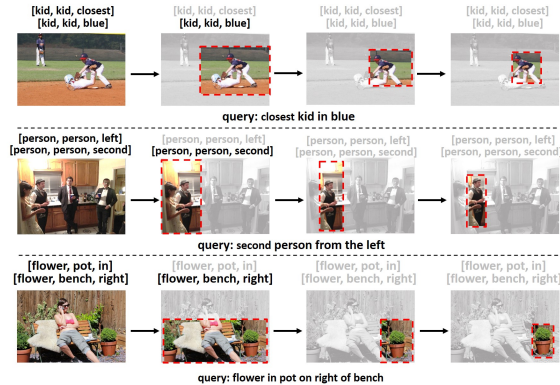


Figure 4. Visualization of the reasoning process about how to iteratively localize the target. The image patch generated through shrinking is marked in red box. The parts above the images indicate the triads parsed from the query, with the black units representing the contents remaining in the corresponding image patch, and the gray units indicating that the visual contents of the corresponding distracting or reference objects are already removed.




Figure 5. Some failure cases of the proposed method where the green box is the ground-truth one, and the red box is the predicted one which is incorrect.

and Table 3. However, the accuracy gains on RefCOCO+ and RefClef are not as large as that on RefCOCOg. We believe the reason is that the proportion of “targets with reference” in RefCOCO+, as well as RefClef, is relatively smaller than that in RefCOCOg. Ultimately, when it comes to the RefCOCO dataset, where most queries are simple and “targets with reference” are rare, the proposed method only achieves a comparable accuracy, as the superiority of better utilizing the contextual relation information turns useless.

Aforementioned evaluation results on different datasets further prove the positive correlation between the accuracy gain of the proposed method and the number of “targets with reference”. Though we have not achieved accuracy gain on all datasets, the average accuracy obtained by the proposed method, on the RefCOCOg, RefCOCO+ and RefCOCO datasets, is still higher than the previous SOTA method RCCF [22] by 1.33%.

Apart from the accuracy gain, it can also be observed from Table 2 that the performance of proposed method is more stable and robust on different datasets compared with other methods. For instance, RCCF [22] achieves quite high accuracy (78.93%) on testA of the RefCOCO dataset, but it does not perform well on the RefCOCOg dataset (65.73%), with a gap of more than 13%. The proposed method’s per-

formance is much more stable, with the corresponding accuracy variation less than 5%, demonstrating the robustness of our proposed method.

4.3. Visualization of Reasoning Process

The reasoning process about how to iteratively find and localize the target can be visualized through the sequential shrinking processes, as shown in Fig.4. For instance, in the first case in Fig.4, the first step of the RL agent is to remove the first distracting object, the kid far from us, which is against the query information “closest”, making its corresponding triad [kid,kid,closest] gray. Then, the second step is to eliminate the second distracting object, the kid in white, which is against the query information “blue”, making its corresponding triad [kid,kid,blue] gray. Ultimately, after a few small-stride shrinks, the precise target region can be obtained. Note that in the third case, the triad [flower,pot,in] turns gray because the distracting flowers, which are not in the pot, are removed from the image. In general, the distracting and reference objects are iteratively excluded in the early stage, and then the target region is gradually refined, until the exact bounding box of the target is obtained.

Some failure cases of the proposed methods are also reported in Fig.5. In Fig.5.(a), the heavily-occluded target is difficult to be localized. Besides, tiny facial difference is also hard to be recognized, like the man and the boy wearing the same clothes in Fig.5.(b). The similar challenge is also encountered when dealing with the difference in the textual description about the clothes styles, as shown in Fig.5.(c). We believe the reason for these failure cases is that the similar situations are very rare in the training set, making them quite tough to be handled for REG methods. We will try to address these hard cases in our future works.

4.4. Ablation Study

4.4.1 RL State Design

The first ablation study is to explore the contribution of the linguistic, visual and spatial components in the proposed state, respectively. As shown in Table 4, for the linguistic component, the triad-based feature is replaced with the feature extracted through Bert [8] (**w/o triad**). Then, the multi-scale feature maps are not utilized, with the raw feature extracted from the CNN network directly employed as the final visual feature (**w/o multi-scale**), to investigate the contribution of the multi-scale visual feature. Lastly, the spatial component is removed from the state to check its influence (**w/o spatial**). The proposed method shows an accuracy gain against all aforementioned settings, which demonstrates the significance of each component in the proposed RL state. Note that the accuracy gain against **w/o triad** comes from better feature representation for the query sen-

tence using the triads, as trivial words in the query (e.g., articles, conjunctions) are eliminated, and the key words are organized into a unified format, allowing a better multi-modal comprehension.

4.4.2 RL Action Design

The second ablation study is to study different designs for the agent action. As shown in Table 4, firstly, the adaptive shrinking stride is replaced with the fixed shrinking stride (**fixed stride**). Specifically, within each iteration, the shrinking stride is fixed as αH for a_1 and a_2 , and αW for a_3 and a_4 , where H and W are the width and height of the original image img , respectively, and α is set as 0.2, being the same as in the proposed adaptive shrinking mechanism. The performance drops dramatically in this setting, demonstrating an adaptive shrinking stride contributes significantly to precise localization. Apart from the actions in the shrinking process, we further disable the bounding box refinement network described in Section 3.5 (**w/o refinement**), with the comparison result demonstrating the contribution of the bounding box refinement step.

4.4.3 RL or Supervised Learning?

The third ablation study is to explore whether the proposed method can be trained in a supervised way, and we introduce **supervised training** in Table 4. In this setting, we randomly sample a region, which covers the whole target object box and its size is larger than the target object box. Then, one side of the sampled region, among the top, bottom, left and right sides, with the highest spatial difference against the target region, is set as the positive label indicating the shrinking direction, if the difference is greater than a threshold. Otherwise, the stop action is set as the positive label. As can be observed in Table 4, the proposed RL performs much better than the supervised training. We believe that the major reason is that the adopted sampling strategy for the supervised training cannot cover all situations in reality. In addition, the RL training can consider not only the current profit but also the potential profit in the future, making a higher overall performance.

5. Conclusion

We have proposed an iterative shrinking mechanism to tackle the REG task, where the image iteratively shrinks into the correct target region. An RL agent is adopted to decide the shrinking direction within each iteration, where all objects in the image patch, including the target, reference and distracting objects, are comprehensively considered to predict the optimal shrinking direction. In future, we plan to deceptively eliminate the contextual information whose corresponding image region has already been removed.

References

- [1] Ankan Bansal, Yuting Zhang, and Rama Chellappa. Visual question answering on image sets. *ECCV*, 2020. [4321](#)
- [2] Danqi Chen and Christopher D Manning. A fast and accurate dependency parser using neural networks. In *EMNLP*, 2014. [4324](#), [4326](#)
- [3] Hui Chen, Guiguang Ding, Xudong Liu, Zijia Lin, Ji Liu, and Jungong Han. IMRAM: Iterative matching with recurrent attention memory for cross-modal image-text retrieval. In *CVPR*, 2020. [4321](#)
- [4] Long Chen, Xin Yan, Jun Xiao, Hanwang Zhang, Shiliang Pu, and Yueting Zhuang. Counterfactual samples synthesizing for robust visual question answering. In *CVPR*, 2020. [4321](#)
- [5] Shizhe Chen, Qin Jin, Peng Wang, and Qi Wu. Say as you wish: Fine-grained control of image caption generation with abstract scene graphs. In *CVPR*, 2020. [4321](#)
- [6] Tianlang Chen, Jiajun Deng, and Jiebo Luo. Adaptive offline quintuplet loss for image-text matching. *ECCV*, 2020. [4321](#)
- [7] Marcella Cornia, Matteo Stefanini, Lorenzo Baraldi, and Rita Cucchiara. Meshed-memory transformer for image captioning. In *CVPR*, 2020. [4321](#)
- [8] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. Bert: Pre-training of deep bidirectional transformers for language understanding. *preprint arXiv:1810.04805*, 2018. [4328](#)
- [9] Longteng Guo, Jing Liu, Xinxin Zhu, Peng Yao, Shichen Lu, and Hanqing Lu. Normalized and geometry-aware self-attention network for image captioning. In *CVPR*, 2020. [4321](#)
- [10] Tanmay Gupta, Arash Vahdat, Gal Chechik, Xiaodong Yang, Jan Kautz, and Derek Hoiem. Contrastive learning for weakly supervised phrase grounding. *ECCV*, 2020. [4321](#)
- [11] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *CVPR*, pages 770–778, 2016. [4324](#), [4326](#)
- [12] Sepp Hochreiter and Jürgen Schmidhuber. Long short-term memory. *Neural Computation*, 9(8):1735–1780, 1997. [4322](#)
- [13] Matthew Honnibal and Ines Montani. spacy 2: Natural language understanding with bloom embeddings, convolutional neural networks and incremental parsing. 2017. [4324](#)
- [14] Tommi Jaakkola, Satinder P Singh, and Michael I Jordan. Reinforcement learning algorithm for partially observable markov decision problems. In *NIPS*, 1995. [4322](#), [4323](#)
- [15] Huaizu Jiang, Ishan Misra, Marcus Rohrbach, Erik Learned-Miller, and Xinlei Chen. In defense of grid features for visual question answering. In *CVPR*, 2020. [4321](#)
- [16] Sahar Kazemzadeh, Vicente Ordonez, Mark Matten, and Tamara Berg. Referitgame: Referring to objects in photographs of natural scenes. In *EMNLP*, 2014. [4326](#)
- [17] Junyeong Kim, Minuk Ma, Trung Pham, Kyungsu Kim, and Chang D Yoo. Modality shifting attention network for multi-modal video question answering. In *CVPR*, 2020. [4321](#)
- [18] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *preprint arXiv:1412.6980*, 2014. [4326](#)
- [19] Vijay R Konda and John N Tsitsiklis. Actor-critic algorithms. In *NIPS*, 2000. [4323](#), [4325](#)
- [20] Thao Minh Le, Vuong Le, Svetha Venkatesh, and Truyen Tran. Hierarchical conditional relation networks for video question answering. In *CVPR*, 2020. [4321](#)
- [21] Kungpeng Li, Yulun Zhang, Kai Li, Yuanyuan Li, and Yun Fu. Visual semantic reasoning for image-text matching. In *CVPR*, 2019. [4321](#)
- [22] Yue Liao, Si Liu, Guanbin Li, Fei Wang, Yanjie Chen, Chen Qian, and Bo Li. A real-time cross-modality correlation filtering method for referring expression comprehension. In *CVPR*, 2020. [4322](#), [4323](#), [4326](#), [4327](#)
- [23] Tsung-Yi Lin, Piotr Dollár, Ross Girshick, Kaiming He, Bharath Hariharan, and Serge Belongie. Feature pyramid networks for object detection. In *CVPR*, 2017. [4324](#), [4326](#)
- [24] Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. Microsoft coco: Common objects in context. In *ECCV*, 2014. [4326](#)
- [25] Chunxiao Liu, Zhendong Mao, Tianzhu Zhang, Hongtao Xie, Bin Wang, and Yongdong Zhang. Graph structured network for image-text matching. In *CVPR*, 2020. [4321](#)
- [26] Daqing Liu, Hanwang Zhang, Feng Wu, and Zheng-Jun Zha. Learning to assemble neural module tree networks for visual grounding. In *CVPR*, 2019. [4321](#)
- [27] Xuejing Liu, Liang Li, Shuhui Wang, Zheng-Jun Zha, Dechao Meng, and Qingming Huang. Adaptive reconstruction network for weakly supervised referring expression grounding. In *ICCV*, 2019. [4321](#)
- [28] Xihui Liu, Zihao Wang, Jing Shao, Xiaogang Wang, and Hongsheng Li. Improving referring expression grounding with cross-modal attention-guided erasing. In *CVPR*, 2019. [4323](#)
- [29] Gen Luo, Yiyi Zhou, Xiaoshuai Sun, Liujuan Cao, Chenglin Wu, Cheng Deng, and Rongrong Ji. Multi-task collaborative network for joint referring expression comprehension and segmentation. In *CVPR*, 2020. [4326](#), [4327](#)
- [30] Junhua Mao, Jonathan Huang, Alexander Toshev, Oana Camburu, Alan L Yuille, and Kevin Murphy. Generation and comprehension of unambiguous object descriptions. In *CVPR*, 2016. [4322](#), [4326](#)
- [31] Tomas Mikolov, Kai Chen, Greg Corrado, and Jeffrey Dean. Efficient estimation of word representations in vector space. *preprint arXiv:1301.3781*, 2013. [4324](#)
- [32] Volodymyr Mnih, Adria Puigdomenech Badia, Mehdi Mirza, Alex Graves, Timothy Lillicrap, Tim Harley, David Silver, and Koray Kavukcuoglu. Asynchronous methods for deep reinforcement learning. In *ICML*, 2016. [4325](#)
- [33] Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Alex Graves, Ioannis Antonoglou, Daan Wierstra, and Martin Riedmiller. Playing atari with deep reinforcement learning. *arXiv:1312.5602*, 2013. [4323](#)
- [34] Yingwei Pan, Ting Yao, Yehao Li, and Tao Mei. X-linear attention networks for image captioning. In *CVPR*, 2020. [4321](#)
- [35] Jeffrey Pennington, Richard Socher, and Christopher D Manning. Glove: Global vectors for word representation. In *EMNLP*, 2014. [4324](#), [4326](#)

- [36] Bryan A Plummer, Paige Kordas, M Hadi Kiapour, Shuai Zheng, Robinson Piramuthu, and Svetlana Lazebnik. Conditional image-text embedding networks. In *ECCV*, 2018. [4326](#)
- [37] Joseph Redmon and Ali Farhadi. Yolov3: An incremental improvement. *preprint arXiv:1804.02767*, 2018. [4322](#)
- [38] Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun. Faster r-cnn: Towards real-time object detection with region proposal networks. In *NIPS*, 2015. [4322](#)
- [39] Steven J Rennie, Etienne Marcheret, Youssef Mroueh, Jerret Ross, and Vaibhava Goel. Self-critical sequence training for image captioning. In *CVPR*, 2017. [4323](#)
- [40] Anna Rohrbach, Marcus Rohrbach, Ronghang Hu, Trevor Darrell, and Bernt Schiele. Grounding of textual phrases in images by reconstruction. In *ECCV*, 2016. [4322](#)
- [41] Arka Sadhu, Kan Chen, and Ram Nevatia. Zero-shot grounding of objects from natural language queries. In *CVPR*, 2020. [4322](#)
- [42] Fawaz Sammani and Luke Melas-Kyriazi. Show, edit and tell: A framework for editing image captions. In *CVPR*, 2020. [4321](#)
- [43] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. *preprint arXiv:1409.1556*, 2014. [4324](#)
- [44] Mingjie Sun, Jimin Xiao, Eng Gee Lim, Si Liu, and John Yannis Goulermas. Discriminative triad matching and reconstruction for weakly referring expression grounding. *TPAMI*, 2021. [4324](#)
- [45] Mingjie Sun, Jimin Xiao, Eng Gee Lim, Yanchun Xie, and Jiashi Feng. Adaptive roi generation for video object segmentation using reinforcement learning. *Pattern Recognition*, 106:107465, 2020. [4323](#)
- [46] Mingjie Sun, Jimin Xiao, Eng Gee Lim, Bingfeng Zhang, and Yao Zhao. Fast template matching and update for video object tracking and segmentation. In *CVPR*, 2020. [4323](#)
- [47] Richard S Sutton, David A McAllester, Satinder P Singh, and Yishay Mansour. Policy gradient methods for reinforcement learning with function approximation. In *NIPS*, 2000. [4323](#)
- [48] Ruixue Tang, Chao Ma, Wei Emma Zhang, Qi Wu, and Xiaokang Yang. Semantic equivalent adversarial data augmentation for visual question answering. *ECCV*, 2020. [4321](#)
- [49] Alasdair Tran, Alexander Mathews, and Lexing Xie. Transform and tell: Entity-aware news image captioning. In *CVPR*, 2020. [4321](#)
- [50] Ramakrishna Vedantam, C Lawrence Zitnick, and Devi Parikh. Cider: Consensus-based image description evaluation. In *CVPR*, 2015. [4323](#)
- [51] Haoran Wang, Ying Zhang, Zhong Ji, Yanwei Pang, and Lin Ma. Consensus-aware visual-semantic embedding for image-text matching. *ECCV*, 2020. [4321](#)
- [52] Liwei Wang, Yin Li, and Svetlana Lazebnik. Learning deep structure-preserving image-text embeddings. In *CVPR*, 2016. [4322](#)
- [53] Peng Wang, Qi Wu, Jiewei Cao, Chunhua Shen, Lianli Gao, and Anton van den Hengel. Neighbourhood Watch: Referring expression comprehension via language-guided graph attention networks. In *CVPR*, 2019. [4321](#)
- [54] Xinyu Wang, Yuliang Liu, Chunhua Shen, Chun Chet Ng, Canjie Luo, Lianwen Jin, Chee Seng Chan, Anton van den Hengel, and Liangwei Wang. On the general value of evidence, and bilingual scene-text visual question answering. In *CVPR*, 2020. [4321](#)
- [55] Kelvin Xu, Jimmy Ba, Ryan Kiros, Kyunghyun Cho, Aaron Courville, Ruslan Salakhudinov, Rich Zemel, and Yoshua Bengio. Show, attend and tell: Neural image caption generation with visual attention. In *ICML*, 2015. [4323](#)
- [56] Sibe Yang, Guanbin Li, and Yizhou Yu. Cross-modal relationship inference for grounding referring expressions. In *CVPR*, 2019. [4321](#)
- [57] Sibe Yang, Guanbin Li, and Yizhou Yu. Dynamic graph attention for referring expression comprehension. In *ICCV*, 2019. [4326](#), [4327](#)
- [58] Zhengyuan Yang, Tianlang Chen, Liwei Wang, and Jiebo Luo. Improving one-stage visual grounding by recursive subquery construction. *ECCV*, 2020. [4322](#)
- [59] Zhengyuan Yang, Boqing Gong, Liwei Wang, Wenbing Huang, Dong Yu, and Jiebo Luo. A fast and accurate one-stage approach to visual grounding. In *CVPR*, 2019. [4322](#), [4323](#), [4326](#), [4327](#)
- [60] Raymond Yeh, Jinjun Xiong, Wen-Mei Hwu, Minh Do, and Alexander Schwing. Interpretable and globally optimal prediction for textual grounding using image concepts. In *NIPS*, 2017. [4326](#)
- [61] Licheng Yu, Zhe Lin, Xiaohui Shen, Jimei Yang, Xin Lu, Mohit Bansal, and Tamara L Berg. Mattnet: Modular attention network for referring expression comprehension. In *CVPR*, 2018. [4321](#), [4322](#), [4326](#), [4327](#)
- [62] Licheng Yu, Patrick Poirson, Shan Yang, Alexander C Berg, and Tamara L Berg. Modeling context in referring expressions. In *ECCV*, 2016. [4322](#), [4326](#)
- [63] Licheng Yu, Hao Tan, Mohit Bansal, and Tamara L Berg. A joint speaker-listener-reinforcer model for referring expressions. In *CVPR*, 2017. [4321](#), [4323](#), [4326](#), [4327](#)
- [64] Qi Zhang, Zhen Lei, Zhaoxiang Zhang, and Stan Z Li. Context-aware attention network for image-text retrieval. In *CVPR*, 2020. [4321](#)
- [65] Yuanen Zhou, Meng Wang, Daqing Liu, Zhenzhen Hu, and Hanwang Zhang. More grounded image captioning by distilling image-text matching model. In *CVPR*, 2020. [4321](#)