

Modeling Explicit Concerning States for Reinforcement Learning in Visual Dialogue

Zipeng Xu¹

xuzp@bupt.edu.cn

Fandong Meng²

fandongmeng@tencent.com

Xiaojie Wang¹

xjwang@bupt.edu.cn

Duo Zheng¹

zd@bupt.edu.cn

Chenxu Lv¹

chenxulv@bupt.edu.cn

Jie Zhou²

withtomzhou@tencent.com

¹ Beijing University of Posts and Telecommunications,
Beijing, China

² Pattern Recognition Center,
WeChat AI,
Tencent Inc., China

Abstract

To encourage AI agents to conduct meaningful Visual Dialogue (VD), the use of Reinforcement Learning has been proven potential. In Reinforcement Learning, it is crucial to represent states and assign rewards based on the action-caused transitions of states. However, the state representation in previous Visual Dialogue works uses the textual information only and its transitions are implicit. In this paper, we propose Explicit Concerning States (ECS) to represent what visual contents are concerned at each round and what have been concerned throughout the Visual Dialogue. ECS is modeled from multimodal information and is represented explicitly. Based on ECS, we formulate two intuitive and interpretable rewards to encourage the Visual Dialogue agents to converse on diverse and informative visual information. Experimental results on the VisDial v1.0 dataset show our method enables the Visual Dialogue agents to generate more visual coherent, less repetitive and more visual informative dialogues compared with previous methods, according to multiple automatic metrics, human study and qualitative analysis.¹

1 Introduction

Visual Dialogue (VD) aims to create intelligent conversational systems that understand vision and language. It is one of the ultimate goals in Artificial Intelligence (AI) [27]. Although important, it is still challenging to generate coherent, visual-related and informative utterances based on the complete understanding of vision and dialog history.

To approach the target, Das *et al.* [8] propose a Q-Bot-A-Bot image-guessing game, where Q-Bot (who cannot see the image) raises question at each round and A-Bot (who can

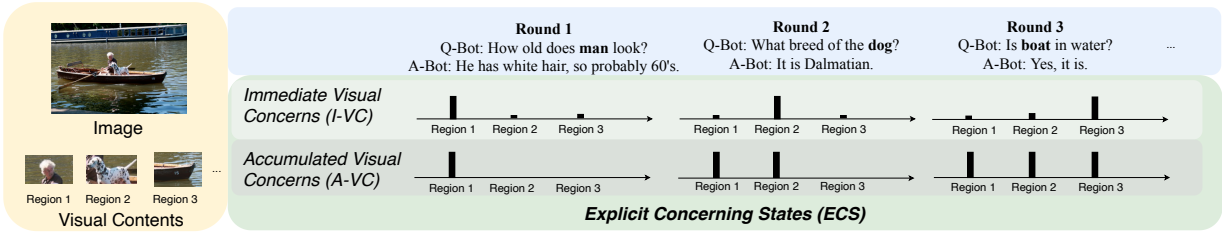


Figure 1: Illustration of Explicit Concerning State (ECS) in the Visual Dialogue process.

see the image) answers, and introduce the use of Reinforcement Learning (RL) to optimize the two Visual Dialogue agents goal-orientedly. In RL, a suitable state representation is regarded as a fundamental part in the learning process [12] and has long been studied in textual dialogue tasks [11, 17, 25, 28, 31]. Similarly in Visual Dialogue, it is vital that the state representation can be able to help distinguish whether the currently generated utterances bring about new useful visual information. Consequently, the reward, which is formulated from the transitions of states, can constantly encourage the agents to converse on new visual informative utterances and achieve success in image-guessing.

However, previous works [3, 8, 24] represent the state in Visual Dialogue in an insufficient and implicit way, resulting in ineffective rewards. As they derive the state representation in Visual Dialogue from the hidden state of the RNN-encoded textual dialogue history, there are several disadvantages: 1) it is insufficient to represent the Visual Dialogue state using textual information merely; 2) the transitions of states are implicit for it can only be measured between successive hidden states; 3) as dialogue history is comparatively long, the change of hidden state before and after the encoding of new question-answer pair can be subtle. Since the rewards judge whether the new utterances are informative based on the distances between the state representation and target image at successive rounds, thus its effects diminish as dialogue progresses. Experimental results show that the agents always repeat themselves, and the image-guessing performance stops increasing at an early stage, indicating the rewards are unable to constantly encourage agents to generate new visual informative utterances. Vishvak *et al.* [24] try to fix dialogue repetition by penalizing the similarity of dialogue hidden states between successive rounds. Yet, this method cannot technically avoid repetition between every alternate dialogue round (at t and $t+2$). Besides, the auxiliary objective built from single modality is inadequate for the multi-modal task.

To remedy the problems, in this paper, we propose Explicit Concerning States (ECS) to represent the state in Visual Dialogue efficiently and explicitly. ECS is a 2-tuple: \langle Immediate Visual Concerns (I-VC), Accumulated Visual Concerns (A-VC) \rangle . As in Fig. 1, I-VC models the visual contents that Q-Bot concerns immediately, A-VC models the visual contents that Q-Bot has concerned throughout the dialogue. I-VC is formalized as the question-history guided visual attention distribution on target image. A-VC is the accumulated binary states, representing whether the corresponding contents have been covered or not. Contrary to previous methods, the advantages of ECS are as follows: 1) ECS is efficiently established based on both visual and textual information; 2) the changing Visual Dialogue states are explicitly related to the visual contents in image; 3) whether Q-Bot refers new visual contents related to the target image can be clearly measured by A-VC.

Based on ECS, we further propose two intuitive and interpretable rewards to encourage the agents to constantly generate new, visual related and informative utterances. Diversity Reward encourages the diversity of I-VC at successive rounds so as to encourage Q-Bot

to ask different visual information. Informativity Reward encourages the growth of A-VC so as to encourage Q-Bot to constantly ask new visual information as dialogue progresses. Meanwhile, A-Bot is also improved as correct and detailed response to current question will guide Q-Bot to transfer its attention to new visual contents.

We evaluate our approach on the large-scale VisDial v1.0 dataset [6]. Experimental results demonstrate the effectiveness and superiority of our approach compared with previous methods, according to multiple automatic metrics, human study and qualitative analysis. To conclude, our main contributions are as follows:

- We propose Explicit Concerning States (ECS) to explicitly model what is currently concerned and what has been concerned in the Visual Dialogue process.
- We propose two intuitive and interpretable rewards based on Explicit Concerning States (ECS) to award the agents for they generate diverse and new visual informative utterances in Reinforcement Learning.
- We conduct experiments on the large-scale VisDial v1.0 dataset and achieve not only improved image-guessing performance, but also less repetitive and more diverse Q-Bot, descriptive and detailed A-Bot. Human study shows our agents generate more coherent, visual-related and informative dialogues that achieve superior image retrieval results as good as human dialogues.

2 Background

To improve Visual Dialog agents with Reinforcement Learning, previous works introduce a Q-Bot-A-Bot image-guessing game setting. Given an undisclosed image I , Q-Bot, who only knows the image caption c at first, has to raise a series of questions about the undisclosed image and make guesses. A-Bot, who can see the image, gives its answer accordingly.

Q-Bot consists of an encoder to encode the textual dialogue history $H_t = \{c, (q_1, a_1), \dots, (q_{t-1}, a_{t-1})\}$, a decoder to generate question q_t and a feature regression network to make image prediction \hat{y}_t . A-Bot consists of a multi-modal encoder, which encodes the image, dialogue history and question, and a decoder, to generate answer a_t .

Under the setting, RL is used to jointly optimize Q-Bot and A-Bot by assigning rewards based on the action-caused transitions of state. At round t , given dialogue history $H_t = \{c, (q_1, a_1), \dots, (q_{t-1}, a_{t-1})\}$, Q-Bot generates q_t and A-Bot generates a_t . In previous methods [3, 8, 24], w.r.t $S_t = \{I, H_t\}$ and action A_t (i.e., selecting the words from the vocabulary V to generate q_t and a_t), the reward is formulated as:

$$r(S_t, A_t) = \|\mathbf{y}^{\text{gt}} - \hat{\mathbf{y}}_{t-1}\|_2^2 - \|\mathbf{y}^{\text{gt}} - \hat{\mathbf{y}}_t\|_2^2, \quad (1)$$

where \mathbf{y}^{gt} is the image feature of target I and $\hat{\mathbf{y}}_t$ is the predicted image feature, which is derived from Q-Bot by passing the hidden state of the RNN-encoded textual dialogue through the feature regression network, i.e., a fully-connected layer. REINFORCE algorithm [26] is used to update agents' parameters with the designed reward.

3 Explicit Concerning States in VD

Explicit Concerning States (ECS) is a 2-tuple: <Immediate Visual Concerns (I-VC), Accumulated Visual Concerns (A-VC)>. As in Fig. 2, I-VC models what visual contents that

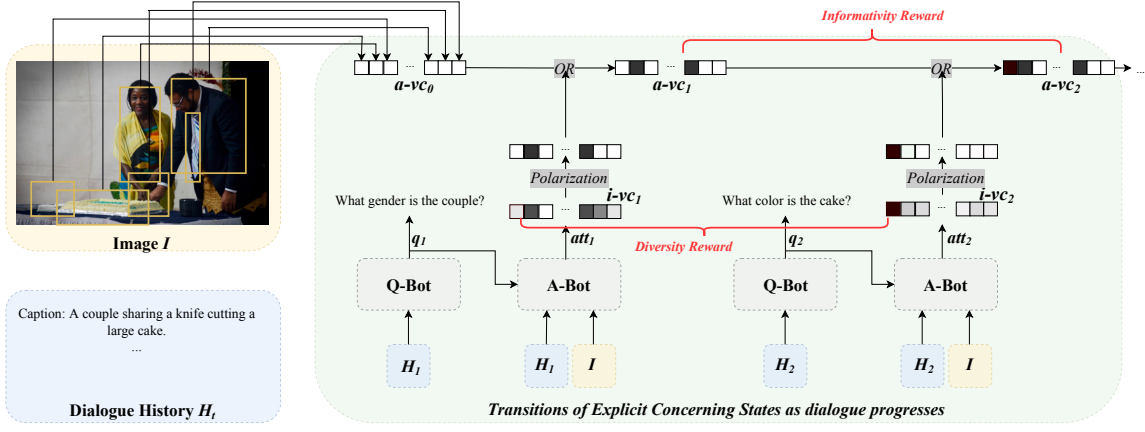


Figure 2: Overview of the proposed ECS and ECS-based reward formulations.

Q-Bot concerns immediately, and A-VC models what visual contents that Q-Bot has concerned throughout the dialogue. As Q-Bot and A-Bot are jointly optimized in a cooperative setting, ECS is modeled from A-Bot’s view, concretely based on some existed computation in an attention-based A-Bot.

Firstly, we propose I-VC to model the immediately concerned visual contents. Considering the co-reference phenomenon, we formalize I-VC based on the question-history-guided attention on image in an attention-based A-Bot. As the common visual attention mechanisms in A-Bot, at round t , current question correlated visual attention is guided by current question q_t and history H_t . Accordingly, I-VC at round t is formulated as:

$$\mathbf{i}\text{-vc}_t = \mathbf{att}_t = P_{A\text{-Bot}}(o|I, H_t, q_t), \quad (2)$$

where o is the set of all objects or spatial grids in I and $i\text{-vc}_t^j$ is the attention weight for $o_j \in o$.

Further, we propose A-VC to model the accumulated concerned visual contents, i.e., the visual contents that have been concerned by Q-Bot throughout the dialogue. A-VC is formulated as the accumulated binary states correspond with $\{o_j\}_{j=1}^N$, where N is the number of all object proposals or spatial grids. At t -th round, $a\text{-vc}_t^j = 0$ means the o_j has not been referred in the past t -round dialogue while $a\text{-vc}_t^j = 1$ otherwise. In concrete, A-VC is initialized with $\mathbf{0} \in R^N$, then is updated by two steps as dialogue progresses:

Step 1, compute the polarized I-VC by projecting the attention weights in I-VC into binary states by the following operations:

$$\text{norm}(\mathbf{i}\text{-vc}_t) = \frac{\mathbf{i}\text{-vc}_t - \min(\mathbf{i}\text{-vc}_t)}{\max(\mathbf{i}\text{-vc}_t) - \min(\mathbf{i}\text{-vc}_t)}, \quad (3)$$

$$\text{polarize}(i\text{-vc}_t^j) = \begin{cases} 1, & \text{if } \text{norm}(i\text{-vc}_t^j) > \gamma, \\ 0, & \text{else.} \end{cases} \quad (4)$$

Accordingly, $i\text{-vc}_t^j \in (0, 1)$ is projected into a binary value in $\{0, 1\}$ with a threshold γ . This is conducted to have a definite identification of each visual content $o_j \in o$ is immediately concerned by Q-Bot or not: the $\text{polarize}(i\text{-vc}_t^j) = 1$ means Q-Bot’s immediately concerned visual contents include o_j while $\text{polarize}(i\text{-vc}_t^j) = 0$ otherwise.

Step 2, update A-VC with the polarized I-VC by the logic OR operation:

$$\mathbf{a}\text{-vc}_t = \text{OR}(\mathbf{a}\text{-vc}_{t-1}, \text{polarize}(\mathbf{i}\text{-vc}_t)) \quad (5)$$

At each round, A-VC is dynamically updated. Only when current question concerns new visual contents (e.g., o_j , whose $a-vc_{t-1}^j$ is 0 and $polarize(i-vc_t^j)$ is 1), will it lead to the growth of $\mathbf{a-vc}_t$, because $a-vc_t^j$ will turn from 0 to 1. Accordingly, whether Q-Bot refers to new visual contents in the image is clearly measured by comparing A-VC at different rounds.

With I-VC and A-VC, ECS efficiently and explicitly represents the state in Visual Dialogue. Using both visual and textual information enables more accurate modeling results as compared with previous single modal methods. I-VC and A-VC thus provide reliable references for reward formulations.

4 ECS-based RL framework

4.1 Learning Environment

Learning Environment includes the supervisedly pretrained Q-Bot and A-Bot. We use the same Q-Bot as in previous work [8]. To establish ECS, we need an attention-based A-Bot. Among many potential models, we choose to use a classic attention-based A-Bot with History-Conditioned Image Attentive Encoder (HCIAE) [16] to verify the effectiveness of our method. For image feature, we use the static features provided by bottom-up attention [5], which is trained based on the framework of Faster-RCNN [20] and the annotated data in Visual Genome [13]. The method adaptively extracts K (ranging from 10 to 100) object proposals for each image. Accordingly, I is represented by $\{\mathbf{i}_1, \mathbf{i}_2, \dots, \mathbf{i}_K\}$.

4.2 Reward Formulations

Based on ECS, we formulate two intuitive and interpretable rewards, which are Diversity Reward (DR) and Informativity Reward (IR), to award agents with reliable reasons at different Visual Dialogue states. Fig. 2 illustrates the two rewards in the Visual Dialogue process.

Diversity Reward (DR): DR is to award the agents for Q-Bot holds different visual concerns at successive rounds. As I-VC represents what visual contents that Q-Bot concerns immediately, DR is formalized as the KL-divergence between successive I-VCs:

$$r^D(S_t, A_t) = D_{KL}(\mathbf{i-vc}_t || \mathbf{i-vc}_{t-1}). \quad (6)$$

On the one hand, DR encourages Q-Bot to successively ask diverse visual contents. On the other hand, it encourages A-Bot to give correct and detailed answers, as they will enable Q-Bot to transfer its attention to new visual contents.

Informativity Reward (IR): IR is to award the agents for Q-Bot constantly concerns new visual contents in the undisclosed image as dialogue progresses. As A-VC represents what visual contents have been covered by Q-Bot throughout the dialogue, IR corresponds with constant growth of A-VC. Concretely, IR is formalized as:

$$r^I(S_t, A_t) = \begin{cases} 1, & \text{if } \sum_{j=1}^N a-vc_t^j - a-vc_{t-1}^j > 0, \\ 0, & \text{else.} \end{cases} \quad (7)$$

Only when Q-Bot asks about new visual contents in current round will the agents be awarded.

To conclude, DR encourages Q-Bot to question about different visual contents at successive rounds: ‘A’ now, ‘B’ next. However, DR can not avoid the repetition at alternate rounds

as it will indiscriminately award ‘A’-‘B’-‘A’-‘B’ and ‘A’-‘B’-‘C’-‘D’. To make up for this, IR is designed to only award the agents for Q-Bot questions new. Jointly using DR and IR enables 1) a Q-Bot that will not repeat itself and will keep exploring the unknown areas, and 2) an A-Bot that tends to give detailed responses so as to guide Q-Bot to transfer its attention to new contents, and consequently the visual informative dialogues.

4.3 Training in RL

Together with the Diversity Reward, the Informativity Reward and the original reward (annotated as r^O) as described in Eq. 1, the final reward is assigned as:

$$r_t = \alpha_O r_t^O + \alpha_D r_t^D + \alpha_I r_t^I, \quad (8)$$

where α_O , α_D and α_I are coefficients for each reward, respectively. With the final reward r_t , we use the same REINFORCE algorithm as previous methods to update policy parameters. More details about the algorithm are given in Appendix A.

5 Experiments

We conduct experiments on the large-scale VisDial v1.0 dataset. The train split consists of 123k dialogues on MS-COCO [15] images and the validation split consists of 2k dialogues on Flickr images. All dialogues last for 10 rounds for each image.

We compare our ECS-based rewards enhanced agents with previous state-of-the-art methods. The comparing methods are: 1) ‘SL’: the supervisedly pretrained basic Q-Bot and A-Bot; 2) ‘RL’: basic Q-Bot and A-Bot fine-tuned by RL using the original reward [8]; 3) ‘RL - diverse’ [24]: diverse Q-Bot (optimized with an additional objective) and A-Bot fine-tuned by RL using the original reward. Our full method is represented as ‘RL-Ours’, i.e., basic Q-Bot and A-Bot fine-tuned by RL using the ECS-based rewards in addition. To make a fair comparison, all methods are implemented based on the same learning environment (as introduced in Sec. 4.1). Training details are given in Appendix B.

5.1 Results

Q-Bot Performance. We evaluate Q-Bot performance with the following metrics: 1) Unique Question: the number of unique questions per dialogue instance; 2) Mutual Overlap [10]: the average value of BLEU-4 overlap of each question in the generated 10-round dialogue with the other 9 questions; 3) Dist-n and Ent-n [14, 33]: the number and entropy of unique n-grams in the generated questions, normalized by the number of total tokens. As shown in Fig. 3(a), jointly using the ECS-based rewards significantly improves Q-Bot performance as compared to previous methods. Our method, i.e., ‘RL-Ours’, generates 8.39 unique question in the 10-round dialogue in average, indicating the agent’s superior ability to conduct less repetitive dialog. Besides, the results on Mutual Overlap, Dist-n and Ent-n show our method helps generate question with higher language diversity.

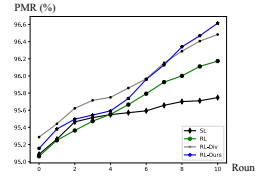
A-Bot Performance. We evaluate A-Bot’s answers to human questions in VisDial v1.0 dataset with the retrieval metrics introduced in [7]. Fig. 3(b) gives the comparing results of A-Bot performance. According to the automatic evaluation metrics, the improvements brought by our method is not significant. However, our goal is to motivate the answering agents to generate more descriptive, detailed and visual informative responses which

	Unique question \uparrow	Mutual overlap \downarrow	Dist-1 \uparrow	Dist-2 \uparrow	Ent-1 \uparrow	Ent-2 \uparrow
<i>SL</i>	5.82	0.67	0.30	0.37	2.55	2.82
<i>RL</i>	6.15	0.61	0.33	0.40	2.61	2.86
<i>RL-Diverse</i>	7.15	0.54	0.35	0.45	2.73	3.06
<i>RL-Ours</i>	8.39	0.43	0.40	0.54	2.89	3.36

(a) Q-Bot performance

	MRR \uparrow	R@1 \uparrow	R@5 \uparrow	Mean \downarrow	NDCG \uparrow
<i>SL</i>	48.70	39.38	57.77	18.82	55.76
<i>RL</i>	49.17	39.40	59.07	17.60	57.81
<i>RL-Diverse</i>	49.34	39.57	59.09	17.55	57.79
<i>RL-Ours</i>	49.41	39.69	58.99	17.59	57.75

(b) A-Bot performance

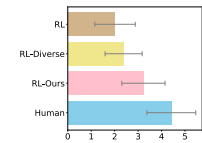
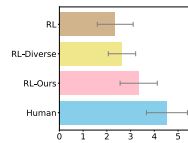
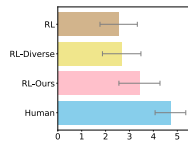
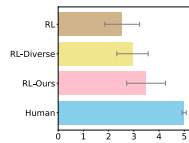


	PMR (%) \uparrow
<i>SL</i>	95.75
<i>RL</i>	96.17
<i>RL-Diverse</i>	96.48
<i>RL-Ours</i>	96.62

(c) Image-guessing performance

Figure 3: Quantitative results on VisDial v1.0 val.

	R@1 \uparrow	R@5 \uparrow
<i>RL</i>	42.00	79.00
<i>RL-Diverse</i>	40.00	79.33
<i>RL-Ours</i>	49.67	91.67
<i>Human</i>	63.33	91.00



(a) Image-Guessing

(b) Coherence

(c) Relevance

(d) Informativity

(e) Detail

Figure 4: Results on human study.

the retrieval-based metrics may not be good at measuring. In the following sections and Appendix C, we further present the distinct characteristics of our A-Bot, which achieves positive results that meet our expectations.

Image-Guessing Performance. Following previous methods, we illustrate the image-guessing performance of each model in Fig. 3(c). Concretely, Q-Bot and A-Bot generate 10-round dialogues via self-talk and each Q-Bot makes the image feature predictions based on the generated dialogues. Then the candidate images in val-set are ranked and PMR (Percentile Mean Rank) is derived accordingly. It is worth mentioning that evaluation standards vary because different prediction models are used. High PMR at this point does not guarantee the high quality of the generated dialog, but its trends indicate whether the rewards in each RL method constantly take effects as dialogue proceeds.

According to the PMR trends across dialogue rounds in the left part of Fig. 3(c): in ‘*SL*’, PMR stops increasing at an early stage. By contrast, Reinforcement Learning takes effects to incentive the growth. Compared to ‘*RL*’, our method presents better performance as the growth on PMR are more significant and persistent, indicating the generated dialogues can constantly provide more effective information for guessing the target image. Compared to previous state-of-the-art method ‘*RL-Diverse*’, which uses extra optimizing objective during the supervise learning stage, our method still achieves better performance.

Human Study. Two human evaluation tasks are conducted: dialogue-based image-guessing task and dialogue quality evaluation task. Six postgraduate students participate in and each evaluates 50 samples for each method variant.

Dialogue-based Image-Guessing task is to retrieve the target image based on the given dialogue. Following previous work [24], evaluators are asked to pick top-1 and top-5 likely images from a 16-candidate-image pool and R@1 and R@5 are derived accordingly. We use the same 16-candidate-image pool which consists of the target image, 5 nearest neigh-


Image + Caption	RL	RL - Diverse	RL - Ours
 <p>Caption: a group of skiers waiting for the lift.</p>	<p>Is the photo in color ?yes is it a professional photo ?yes is it well lit ?yes is it daytime ?yes is it sunny ?no can you see the sky ?yes can you see the clouds ?no is it sunny ?no anything else in the buildings ?no anything else in the buildings ?no</p>	<p>how many people are there ?? are they men or women ?i ca n't tell are they wearing hats ?yes what color are the hats ?yellow , red , black can you see the sky ?yes are there trees ?no can you see buildings around ?no can you see trees around ?no can you see buildings around ?no can you see trees around ?no</p>	<p>how many skiers are there ?6 are they wearing helmets ?yes are they wearing goggles ?yes do they have hats on ?yes is there a lot of snow ?yes do you see any trees ?no do you see any other equipment ?no what color are the shoes ?ca n't tell , they are covered in snow gear is there anything else in the background ?no , just skis and ski poles what color are the shoes ?ca n't tell , they are covered in snow gear</p>

Figure 5: Generated dialogues from comparing methods.

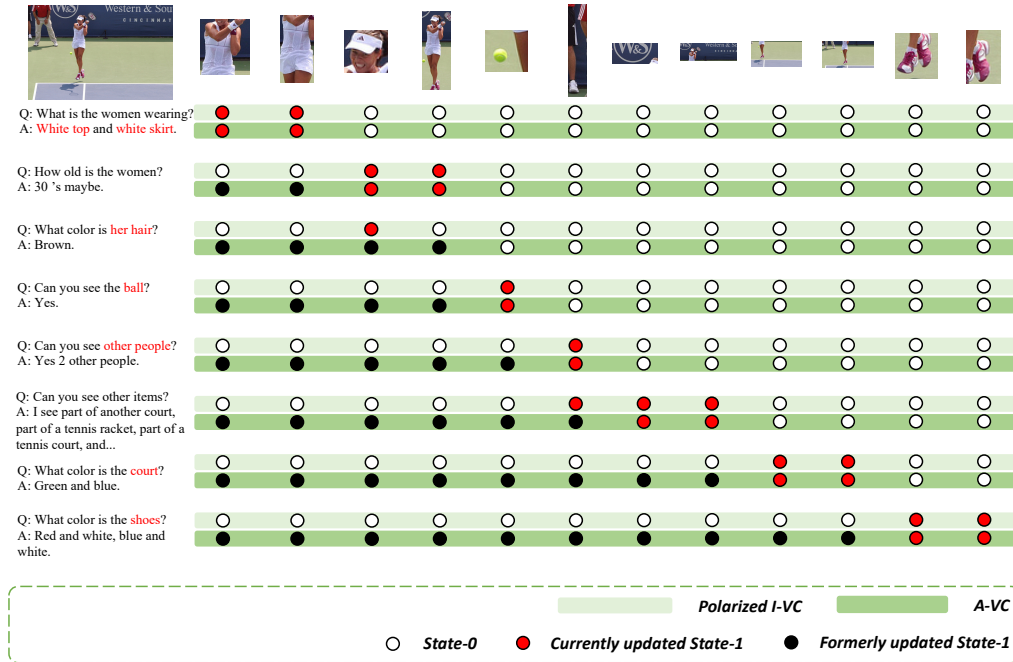


Figure 6: Visualization of ECS.

bors, and 10 random images from the VisDial v1.0 test-std split. The image-guessing performances present whether the generated dialogues are human-understandable, and include enough information that helps to figure out the likely images. As shown in Fig. 4(a), our method achieves best performances on R@1 and R@5 as compared to previous methods. We even outperform human on R@5, indicating that dialogues generated by our methods can cover mostly related visual contents in the target image.

Dialogue Quality Evaluation task is to rate the dialogue quality from 1 to 5 on four criteria: coherence, relevance, informativity and detail. Coherence considers whether the dialogue contexts are natural, coherent, less repetitive and consistent. Relevance assesses the degree that dialogues are relevant to the target image and given caption. Informativity defines how well the dialogues provide adequate information about the target image. Detail presents the extent that dialogues include detailed visual information. As can be seen in Fig. 4(b)-4(e), our method achieves best scores on all criteria as compared to previous methods.

Qualitative Analysis. In Fig. 5, we illustrate the generated dialogue examples from comparing models. Dialogue generated by ‘RL’ is not very related to the main contents of the image. ‘RL-Diverse’ is of more diversity but always repeats at alternate turns. Besides, the description on visual contents of both models are coarse. Contrarily, our method generates dialogues that are less repetitive, more visual-coherent and more informative. As shown,

	Q-Bot		A-Bot		Image-Guessing
	Unique Question \uparrow	Mutual Overlap \downarrow	MRR \uparrow	R@1 \uparrow	PMR (%) \uparrow
<i>RL</i>	6.15	0.61	49.17	39.40	96.17
<i>RL+DR</i>	7.22	0.53	49.19	39.42	96.36
<i>RL+IR</i>	4.86	0.67	49.39	39.55	97.35
<i>RL+DR+IR (RL-Ours)</i>	8.39	0.43	49.41	39.69	96.62

Table 1: Comprehensive quantitative results on ablation study.

dialogue generated by our methods gradually covers different contents in the image, which matches the goal of enlarging A-VC well. Besides, our dialogue agents are able to generate informative utterances with detailed visual descriptions. We are the only one to describe the “skiers”, “wearing helmets” and “wearing goggles”, which are all important details of the image. Also, our A-Bot provides more descriptive and informative answers. For example, when answering “what color are the shoes”, our agent gives additional explanations as “can’t tell, they are covered in snow gear”. Answers like this are more human-like and is helpful to keep dialog-participant interested in. More samples are given in Appendix D.

Besides, we visualize ECS <I-VC, A-VC> in a dialogue process in Fig. 6. As dialogue progresses, Q-Bot step-wisely concerns about different visual contents related to image, raises corresponding visual-coherent questions and gradually collects detailed information of the undisclosed image. Besides, the detailed answers given by A-Bot also helps promote the dialogue process. Noticeably, when A-Bot refers “court”, Q-Bot raises the subsequent question about the court: “What color is the court?”. Correspondingly, the immediate concerned contents are explicitly captured by I-VC and the accumulated concerned contents are explicitly captured by A-VC. As highlighted in red, the textual contents in dialogue align well with the states on image. During the process, I-VC varies and A-VC grows at each round. The phenomena are consistent with our design intuitions.

5.2 Ablation Study

To study the effects of individual reward, i.e., DR and IR, we conduct ablation study. Tab. 1 shows the comprehensive quantitative results evaluated by automatic metric on the three aspects (Q-Bot, A-Bot and image-guessing), for each ablative model.

Individually using DR (*RL+DR*) brings about slightly better Q-Bot’ diversity and image-guessing performance as compared with *RL*, e.g. Q-Bot’s Unique Question is improved from 6.15 to 7.22 and PMR is from 96.17 to 96.36. And as IR encourages the informativity, *RL+IR* achieves a highest PMR of 97.35, indicating the generated dialogues are more informative so as to achieve more success in guessing the target image, but the Q-Bot diversity is poor. This happens because the optimized agents are more to achieving the image-guessing goal and thus ignoring the dialogue quality. Specifically, Q-Bot learns to ask “what else?” at an early stage and A-Bot tells it nearly everything about the image. Jointly using DR and IR achieves the most human-readable results (*RL+DR+IR*), Q-Bot’s Unique Question reaches the highest 8.39 and PMR reaches 96.62. To conclude, DR encourages the different Q-Bot’s visual related questions at successive rounds but cannot promise the constant informativity; IR encourages the informativity of the whole dialogue but lacks guidance for generating dialogue that exchanges the information step by step; DR and IR are complementary to each other for achieving the visual coherent, less repetitive and visual informative dialogues.

Fig. 7 shows the comparing qualitative results among *RL*, *RL+DR* and *RL+DR+IR*. As compared to *RL*, dialogue from *RL+DR* converses on more visual related contents like


Image + Caption	RL	RL + DR	RL + DR + IR (RL-Ours)
 <p>Caption: the scene has 2 men playing doubles tennis</p>	<p>is this in color ? yes are they wearing visors ? 1 is wearing a visor are they male or female ? male how old do they look ? 30 's anything else in the room ? no anything else in the room ? no anything else in the room ? no anything else in the room ? no anything else in the room ? no</p>	<p>how old are the men ? they look like they are in their 20s are they wearing visors ? 1 is wearing a visor are they wearing visors ? 1 is wearing a visor what color is the court ? the court is green and the grass is green what else are in the picture ? just the tennis player what color shirt ? orange what else ? in the picture ? just the tennis player what color shirt ? orange what else ? in the picture ? just the tennis player what color shirt ? orange</p>	<p>is the photo in color ? yes is it a professional photo ? no is it well lit ? yes is it daytime ? yes how old are the men ? 20 's what race are the men ? white what else is in the background ? tennis rackets , a tennis court , part of a tennis ball , part of a blue sky what color is the man 's shirt ? orange anything else in the background ? no , just a baseball glove , part of a tennis ball , part of a tennis player , the what color is the surface ? green</p>

Figure 7: Qualitative results on ablation study.

‘visors’, ‘court’ and ‘shirt’ but suffers from repetition at successive rounds. By adding IR, the generated dialogue repeats less and is more informative.

6 Related Work

Our work is mostly related to previous works that improve open-domain VD using RL. Previous works [3, 8, 24] implicitly represent the VD state as the textual dialogue encoding, which can be insufficient for the multimodal task. To encourage diverse questions, Vishvak *et al.* [24] propose to penalize the similarity between successive textual dialogue encodings, but the difference of textual dialogue encodings cannot promise visual informative dialogues. Contrarily, we represent the VD state in an explicit and multimodal way, enabling more reasonable and interpretable reward formulations to encourage visual informative dialogues.

Works [1, 2, 23, 30, 34] that studies using RL to improve goal-oriented VD in the Guess-what?! task setting [9] are also related. Their goal is to retrieve a predefined object in the image and the answer can only be “yes/no/not available”. Designing intermediate rewards to encourage the goal-oriented visual question generation has been explored by [22, 32], but the rewards are assigned based on external information, which needs additional annotated information and is limited to the task setting. By contrast, our reward formulations are based on internal dialogue state and are applicable in improving open-domain VD.

Many works [6, 16, 18, 19, 21, 29] focus on training a VD answering agent by Supervised Learning (SL). Despite high performance has been achieved according to automatic evaluation metrics, the predicted answers tend to be generic [4]. With the intuitive ECS-based rewards, our method facilitates an intrinsically motivated A-Bot that can generate descriptive and informative answers, and can be used to fine-tune the attention-based SL methods.

7 Conclusion

In this paper, we propose Explicit Concerning States (ECS) to represent the state in Visual Dialogue efficiently and explicitly, and to provide references for rewards in Reinforcement Learning. Based on ECS, we formulate Diversity Reward (DR) and Informativity Reward (IR), to award agents under different Visual Dialogue states. Boosted by the methods, our agents generate less repetitive, more visual-coherent and informative dialogues compared with previous methods according to multiple criteria.

8 Acknowledgements

We thank the reviewers for their comments and suggestions. This work was supported in part by the National Key Research and Development Program of China under the grant of 2020YFF0305302, and NSFC (No. 61906018), MoE-CMCC "Artificial Intelligence" Project (No. MCM20190701).

References

- [1] Ehsan Abbasnejad, Qi Wu, Iman Abbasnejad, Javen Shi, and Anton van den Hengel. An active information seeking model for goal-oriented vision-and-language tasks. *CoRR*, abs/1812.06398, 2018. URL <http://arxiv.org/abs/1812.06398>.
- [2] Ehsan Abbasnejad, Qi Wu, Javen Shi, and Anton van den Hengel. What's to know? uncertainty as a guide to asking goal-oriented questions. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 4150–4159, 2018.
- [3] Akshat Agarwal, Swaminathan Gurumurthy, Vasu Sharma, and Katia P. Sycara. Mind your language: Learning visually grounded dialog in a multi-agent setting. *CoRR*, abs/1808.04359, 2018. URL <http://arxiv.org/abs/1808.04359>.
- [4] Shubham Agarwal, Trung Bui, Joon-Young Lee, Ioannis Konstas, and Verena Rieser. History for visual dialog: Do we really need it? In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 8182–8197, Online, July 2020. Association for Computational Linguistics. doi: 10.18653/v1/2020.acl-main.728. URL <https://www.aclweb.org/anthology/2020.acl-main.728>.
- [5] Peter Anderson, Xiaodong He, Chris Buehler, Damien Teney, Mark Johnson, Stephen Gould, and Lei Zhang. Bottom-up and top-down attention for image captioning and visual question answering. In *Proceedings of the IEEE conference on Computer Vision and Pattern Recognition*, pages 6077–6086, 2018.
- [6] Abhishek Das, Satwik Kottur, Khushi Gupta, Avi Singh, Deshraj Yadav, José MF Moura, Devi Parikh, and Dhruv Batra. Visual dialog. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 326–335, 2017.
- [7] Abhishek Das, Satwik Kottur, Khushi Gupta, Avi Singh, Deshraj Yadav, José M.F. Moura, Devi Parikh, and Dhruv Batra. Visual Dialog. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2017.
- [8] Abhishek Das, Satwik Kottur, José MF Moura, Stefan Lee, and Dhruv Batra. Learning cooperative visual dialog agents with deep reinforcement learning. In *Proceedings of the IEEE international conference on computer vision*, pages 2951–2960, 2017.
- [9] Harm De Vries, Florian Strub, Sarath Chandar, Olivier Pietquin, Hugo Larochelle, and Aaron Courville. Guesswhat?! visual object discovery through multi-modal dialogue. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 5503–5512, 2017.

- [10] Aditya Deshpande, Jyoti Aneja, Liwei Wang, Alexander G. Schwing, and David A. Forsyth. Diverse and controllable image captioning with part-of-speech guidance. *CoRR*, abs/1805.12589, 2018. URL <http://arxiv.org/abs/1805.12589>.
- [11] Matthew Henderson, Blaise Thomson, and Steve Young. Word-based dialog state tracking with recurrent neural networks. In *Proceedings of the 15th Annual Meeting of the Special Interest Group on Discourse and Dialogue (SIGDIAL)*, pages 292–299, 2014.
- [12] Matt Jones and Fabián Canas. Integrating reinforcement learning with models of representation learning. In *Proceedings of the Annual Meeting of the Cognitive Science Society*, 2010.
- [13] Ranjay Krishna, Yuke Zhu, Oliver Groth, Justin Johnson, Kenji Hata, Joshua Kravitz, Stephanie Chen, Yannis Kalantidis, Li-Jia Li, David A. Shamma, Michael S. Bernstein, and Fei-Fei Li. Visual genome: Connecting language and vision using crowdsourced dense image annotations, 2016.
- [14] Jiwei Li, Michel Galley, Chris Brockett, Jianfeng Gao, and Bill Dolan. A diversity-promoting objective function for neural conversation models. In *Proceedings of the 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 110–119, San Diego, California, June 2016. Association for Computational Linguistics. doi: 10.18653/v1/N16-1014. URL <https://www.aclweb.org/anthology/N16-1014>.
- [15] 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 *European conference on computer vision*, pages 740–755. Springer, 2014.
- [16] Jiasen Lu, Anitha Kannan, , Jianwei Yang, Devi Parikh, and Dhruv Batra. Best of both worlds: Transferring knowledge from discriminative learning to a generative visual dialog model. *NIPS*, 2017.
- [17] Nikola Mrkšić, Diarmuid Ó Séaghdha, Tsung-Hsien Wen, Blaise Thomson, and Steve Young. Neural belief tracker: Data-driven dialogue state tracking. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1777–1788, Vancouver, Canada, July 2017. Association for Computational Linguistics. doi: 10.18653/v1/P17-1163. URL <https://www.aclweb.org/anthology/P17-1163>.
- [18] Yulei Niu, Hanwang Zhang, Manli Zhang, Jianhong Zhang, Zhiwu Lu, and Ji-Rong Wen. Recursive visual attention in visual dialog. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 6679–6688, 2019.
- [19] Jiaxin Qi, Yulei Niu, Jianqiang Huang, and Hanwang Zhang. Two causal principles for improving visual dialog. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2020.
- [20] Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun. Faster r-cnn: Towards real-time object detection with region proposal networks. In *Advances in Neural Information Processing Systems*, pages 91–99, 2015.

- [21] Paul Hongsuck Seo, Andreas Lehrmann, Bohyung Han, and Leonid Sigal. Visual reference resolution using attention memory for visual dialog. In *Advances in neural information processing systems*, pages 3719–3729, 2017.
- [22] Pushkar Shukla, Carlos Elmadjian, Richika Sharan, Vivek Kulkarni, Matthew Turk, and William Yang Wang. What should I ask? using conversationally informative rewards for goal-oriented visual dialog. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 6442–6451, Florence, Italy, July 2019. Association for Computational Linguistics. doi: 10.18653/v1/P19-1646. URL <https://www.aclweb.org/anthology/P19-1646>.
- [23] Florian Strub, Harm de Vries, Jérémie Mary, Bilal Piot, Aaron C. Courville, and Olivier Pietquin. End-to-end optimization of goal-driven and visually grounded dialogue systems. In *Joint Conference on Artificial Intelligence*, 2017.
- [24] Murahari Vishvak, Chattopadhyay Prithvijit, Batra Dhruv, Parikh Devi, and Das Abhishek. Improving generative visual dialog by answering diverse questions. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing*, 2019.
- [25] Jason D Williams, Matthew Henderson, Antoine Raux, Blaise Thomson, Alan Black, and Deepak Ramachandran. The dialog state tracking challenge series. *AI Magazine*, 35(4):121–124, 2014.
- [26] Ronald J Williams. Simple statistical gradient-following algorithms for connectionist reinforcement learning. *Machine learning*, 8(3-4):229–256, 1992.
- [27] Terry Winograd. Understanding natural language. *Cognitive psychology*, 3(1):1–191, 1972.
- [28] Chien-Sheng Wu, Andrea Madotto, Ehsan Hosseini-Asl, Caiming Xiong, Richard Socher, and Pascale Fung. Transferable multi-domain state generator for task-oriented dialogue systems. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 808–819, Florence, Italy, July 2019. Association for Computational Linguistics. doi: 10.18653/v1/P19-1078. URL <https://www.aclweb.org/anthology/P19-1078>.
- [29] Jiang Xiaoze, Yu Jing, Sun Yajing, Qin Zengchang, Zhu Zihao, and Wu Qi. Dam: Deliberation, abandon and memory networks for generating detailed and non-repetitive responses in visual dialogue. In *International Joint Conference on Artificial Intelligence (IJCAI)*, 2020.
- [30] Zipeng Xu, Fangxiang Feng, Xiaojie Wang, Yushu Yang, Huixing Jiang, and Zhongyuan Wang. Answer-driven visual state estimator for goal-oriented visual dialogue. In *Proceedings of the 28th ACM International Conference on Multimedia*, MM '20, page 4271–4279, New York, NY, USA, 2020. Association for Computing Machinery. ISBN 9781450379885. doi: 10.1145/3394171.3413668. URL <https://doi.org/10.1145/3394171.3413668>.
- [31] Steve Young, Milica Gašić, Simon Keizer, François Mairesse, Jost Schatzmann, Blaise Thomson, and Kai Yu. The hidden information state model: A practical framework

for pomdp-based spoken dialogue management. *Computer Speech & Language*, 24(2): 150–174, 2010.

- [32] Junjie Zhang, Qi Wu, Chunhua Shen, Jian Zhang, Jianfeng Lu, and Anton van den Hengel. Goal-oriented visual question generation via intermediate rewards. In *Proceedings of the European Conference on Computer Vision*, September 2018.
- [33] Yizhe Zhang, Michel Galley, Jianfeng Gao, Zhe Gan, Xiujun Li, Chris Brockett, and Bill Dolan. Generating informative and diverse conversational responses via adversarial information maximization. In S. Bengio, H. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 31. Curran Associates, Inc., 2018. URL <https://proceedings.neurips.cc/paper/2018/file/23ce1851341ec1fa9e0c259de10bf87c-Paper.pdf>.
- [34] Rui Zhao and Volker Tresp. Learning goal-oriented visual dialog via tempered policy gradient. In *2018 IEEE Spoken Language Technology Workshop (SLT)*, pages 868–875. IEEE, 2018.