

ECOL-R: Encouraging Copying in Novel Object Captioning with Reinforcement Learning

Yufei Wang¹ and Ian D. Wood¹ and Stephen Wan² and Mark Johnson³

Macquarie University, Sydney, Australia¹

CSIRO Data61, Sydney, Australia²

Oracle Digital Assistant, Oracle Corporation³

yufei.wang@students.mq.edu.au, ian.wood@mq.edu.au

stephen.wan@data61.csiro.au


mark.mj.johnson@oracle.com

Abstract

Novel Object Captioning is a zero-shot Image Captioning task requiring describing objects not seen in the training captions, but for which information is available from external object detectors. The key challenge is to select and describe all salient detected novel objects in the input images. In this paper, we focus on this challenge and propose the **ECOL-R** model (Encouraging Copying of Object Labels with Reinforced Learning), a copy-augmented transformer model that is encouraged to accurately describe the novel object labels. This is achieved via a specialised reward function in the SCST reinforcement learning framework (Rennie et al., 2017) that encourages novel object mentions while maintaining the caption quality. We further restrict the SCST training to the images where detected objects are mentioned in reference captions to train the **ECOL-R** model. We additionally improve our copy mechanism via *Abstract Labels*, which transfer knowledge from known to novel object types, and a *Morphological Selector*, which determines the appropriate inflected forms of novel object labels. The resulting model sets new state-of-the-art on the *nocaps* (Agrawal et al., 2019) and *held-out COCO* (Hendricks et al., 2016) benchmarks.

1 Introduction

Novel Object Captioning is a zero-shot Image Captioning task where the captions should mention novel objects (i.e., not seen in the training captions), but for which information is available from external object detectors. To produce high-quality captions, the captioning models should select and describe all salient detected objects and avoid mentioning minor or irrelevant details in the input images. As shown in Figure 1, caption **A** is the best caption among the three because **A** mentions all salient objects in the images without any unnecessary details while **B** mentions *Bread* which is just a



- A** Two Hamburgers served with french fries and drinks.
- B** Two Hamburgers served with bread on the plate.
- C** Two Hamburgers served with french fries.

Figure 1: Caption **A** is the ground-truth caption for the image. Compared with **B** and **C**, **A** is the best caption because it mentions all salient objects (i.e., *Hamburger*, *French Fries* and *Drinks*). We use *Abstract Labels*, that is hypernyms of the objects' detected object labels in the object representations, transferring knowledge from the objects seen in the training captions to novel objects. Our copy mechanism also selects appropriate inflected forms of object labels (i.e., *Hamburgers* vs. *Hamburger*).

minor detail; and **C** misses the salient object *Drink*. This paper aims to develop a captioning model that produces caption **A**.

We use an advanced copy mechanism, similar to the one in Hu et al. (2020), to effectively integrate novel objects. We follow the setup in Agrawal et al. (2019) and use two object detectors: one providing rich object visual features and another providing task specific (including novel) object labels as copy candidates. Our preliminary experiments show that the copy mechanism is infrequently triggered and unable to mention many salient objects in the input images. We propose the **ECOL-R**

model (Encouraging Copying of Object Labels with Reinforced Learning), a copy-augmented transformer model trained in the Self-Critical Sequence Training (SCST) framework (Rennie et al., 2017). SCST with a CIDEr reward (Vedantam et al., 2015) is a standard approach for training the captioning models (Anderson et al., 2018b), but this paper will show that it does not sufficiently encourage the model to use copy operations. We design a new reward function that provides a reward for each copy operation proportional to the caption quality. We further restrict the SCST training to the images that contain at least one word in the ground truth captions that corresponds to one of the detected object labels. With these innovations, the **ECOL-R** model outperforms a SCST baseline and a strong inference encouragement baseline by a large margin.

Our copy mechanism and caption generator incorporate two enhancements to better choose and incorporate novel objects: *a) Abstract Labels* which correspond to hypernyms of the object labels and facilitate knowledge transfer between objects appearing in training captions and novel objects; *b) a Morphological Selector* which determines the correct inflected form of the copied task specific object labels which is similar in purpose to that proposed in (Lu et al., 2018b).

We evaluate the **ECOL-R** model on the novel object captioning benchmark *nocaps* (Agrawal et al., 2019) and *held-out* COCO (Hendricks et al., 2016). The **ECOL-R** model achieves a new state of the art on both benchmarks and generalizes well to in-domain images.

2 Related Work

Popular Image Captioning models include LSTM-based (Anderson et al., 2018b) and Transformer-based decoders (Herdade et al., 2019; Cornia et al., 2020). The visual encoders are often neural object detectors (Anderson et al., 2018b; Wang et al., 2019) producing Region-of-Interest (ROI) vectors. To train the model to copy novel object labels, the Neural Baby Talk model (**NBT**) (Lu et al., 2018a) and follow-up work (Wu et al., 2018; Yao et al., 2017; Li et al., 2019) use copy mechanisms (Vinyals et al., 2015). The copying candidates are labels of salient objects produced by external object detectors. In this paper, we follow previous work by using the Visual Genome object detector from (Anderson et al., 2018b) as the visual

feature extractor and a task specific object detector to provide object labels for copying.

These models are typically trained with the Cross-Entropy loss (CE). This creates a mismatch between the training and testing environments because the evaluation metrics are non-differentiable text-based measures (Ranzato et al., 2015). Self-Critical Sequence Training (SCST) (Rennie et al., 2017) was proposed to address this issue by directly optimizing the inference output using caption-level rewards, such as CIDEr-D (Vedantam et al., 2015).

There are two existing novel object captioning benchmarks: *a) the held-out COCO Benchmark* (Hendricks et al., 2016), constructed by excluding images containing one of eight selected object classes from the standard COCO 2014 benchmark, and *b) nocaps* (Agrawal et al., 2019), which uses the COCO 2017 benchmark for training and provides new validation and test images from the **Open Images** Dataset with over 400 novel objects. Both benchmarks are object-centric and there is no reliable benchmarks that systematically evaluate the quality of generated actions or attributes.

3 Model

Figure 2 provides an overview of the **ECOL-R** model. We refer to the **ECOL-R** model without SCST training as **ECOL**. We describe this model in Sec. 3.1 and our novel reinforced copy encouragement training in Sec. 3.2.

3.1 The ECOL Model

Input Image Objects: Following the setup in Agrawal et al. (2019), we use two object detectors: the Visual Genome object detector from Anderson et al. (2018b), producing image objects and regions G (represented by embedding vectors $[\mathbf{x}_1^g, \dots, \mathbf{x}_{k^g}^g]$) with detailed visual features; and a task specific object detector, producing image objects F (represented by $[\mathbf{x}_1^f, \dots, \mathbf{x}_{k^f}^f]$) and their corresponding labels $L_f = [l_1, \dots, l_{k^f}]$ used as copy candidates during caption generation. We will introduce object representations \mathbf{x}_i below and define them in Eq. 1.

Image Object Representations: Following Anderson et al. (2018b); Lu et al. (2018a), we represent both sets of objects with Region-Of-Interest (ROI, $\mathbf{r}_i \in \mathbb{R}^{2048}$) vectors from the Visual Genome object detector and object positional features ($\mathbf{p}_i \in \mathbb{R}^8$), including bounding box coordinates and size, and an object label confidence score. In addition,

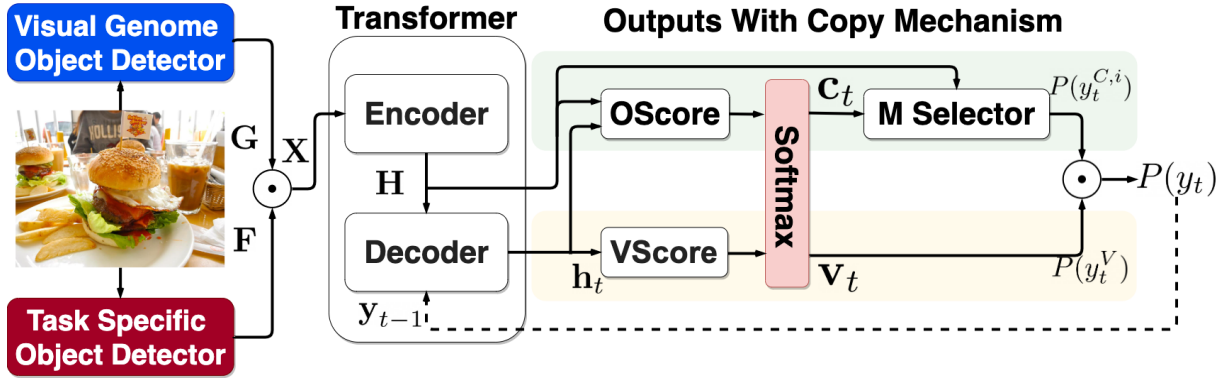


Figure 2: Overview of the **ECOL-R** Model. X is the concatenated Object representations G and F from the two object detectors. The Transformer encoder produces \mathbf{H} and the decoder provides h_t at step t . We then estimate the probabilities for generating each vocabulary word (yellow box) and copying from each task specific image object (green box). The results are concatenated and jointly softmax (red box). We refine each copy probability into the concrete inflected word probability in $MSelector$. The final output $P(y_t)$ concatenates all above word probabilities.

to transfer knowledge from the seen objects to the novel ones, we propose *Abstract Labels* for the task specific objects, described below.

Abstract Labels: The task specific object detectors we use provide taxonomies of object classes, and every detected object is assigned a label from that taxonomy. More general object classes conceptually include all the labels lower in the taxonomy.¹ This provides us with a mechanism for associating class labels not present in the training data with those that do occur in the training data by mapping them to a common ancestor in the hierarchy. Inspired by Ciaramita and Johnson (2003), we define *Abstract Labels* to be a fixed set of ancestor class labels that spans the entire taxonomy (see Figure 3). Using the abstract labels to drive copy decisions allows the usage of known object types to inform the word generation of novel objects. Each object from the task specific detector is associated with its nearest abstract label ancestor. We choose the set of abstract labels such that the objects in the training data are evenly distributed across the set of abstract labels. We represent abstract labels with trainable embeddings $e_i \in \mathbb{R}^d$, where d is the hidden size of our base model. We use the **Open Images V4** class hierarchy for the *nocaps* benchmark and a merged 8 coco super-categories hierarchy Lin et al. (2014) for the *held-out* COCO benchmark. The

¹If the task specific object detector does not provide such a taxonomy, a suitable taxonomy could be obtained from sources such as Wordnet.




Figure 3: A part of the class hierarchy from the Open Images V4 Dataset (Kuznetsova et al., 2018). The green nodes are used as abstract object labels. For each label, its abstract label is its closest green ancestor.

final representation for each object x_i is:

$$x_i = LN(W_r r_i + e_i) + LN(W_p p_i) \quad (1)$$

where LN is layer normalization and $W_r \in \mathbb{R}^{d \times 2048}$, $W_p \in \mathbb{R}^{d \times 8}$ are trainable projections. The two sets of object representations are concatenated as $X = F \odot G$ where \odot represents concatenation.

Transformer Base Model: We use a transformer model (Vaswani et al., 2017) with an N_{enc} -layer encoder and an N_{dec} -layer decoder ($N_{enc} = N_{dec} = 3$ in our experiments). We denote the encoder output $\mathbf{H} = Encoder(X)$. The decoder uses frozen word and positional embeddings WE and PE from GPT2 (Radford et al., 2019) which are helpful in producing better captions describing novel objects. In step t :

$$w_{1:t-1} = WE(y_{1:t-1}) + PE(y_{1:t-1}) \quad (2)$$

$$h_t = Decoder(\mathbf{H}, w_{1:t-1}) \quad (3)$$

where $y_{1:t-1}$ is the generation history and $h_t \in \mathbb{R}^d$.

Outputs With Copy Mechanism: The **ECOL** model either generates words from the vocabulary or copies from task specific objects. We deploy a copy mechanism similar to the dynamic pointer network in Hu et al. (2020). Given the decoder output \mathbf{h}_t , we first calculate a raw score for each vocabulary word:

$$VScore(\mathbf{h}_t) = \mathbf{W}_e \mathbf{h}_t \quad (4)$$

where $\mathbf{W}_e \in \mathbb{R}^{|V| \times d}$, $|V|$ is the GPT2 vocabulary size. We then calculate raw additive attention scores over the encoder output of task specific image objects (i.e., $\mathbf{H}_{1:k^f}$):

$$OScore(\mathbf{H}, \mathbf{h}_t)_i = \mathbf{w}_c^T \tanh(\mathbf{W}_f \mathbf{H}_i + \mathbf{W}_h \mathbf{h}_t) \quad (5)$$

where $i \in [1, k^f]$ and $\mathbf{W}_f \in \mathbb{R}^{d \times d}$, $\mathbf{W}_h \in \mathbb{R}^{d \times d}$ and $\mathbf{w}_c \in \mathbb{R}^d$. Finally, we concatenate the raw scores from $VScore$ and $OScore$ and jointly softmax:

$$[v_t, c_t] = \text{Softmax}([VScore(\mathbf{h}_t) \odot OScore(\mathbf{H}, \mathbf{h}_t)]) \quad (6)$$

where \odot represents concatenation. v_t provides probabilities for GPT2 vocabulary words and c_t provides probabilities for copying task specific object labels.

Morphological Selector: Object labels can appear in inflected forms in captions. For example, in Figure 1, after selecting the object *hamburger*, the **ECOL** model should generate “*hamburgers*” after “*Two*”. We propose a morphological selector (*M Selector*) to refine the copy probability of each task specific image object label l_i (i.e., $c_{t,i}$) into the probabilities of generating all possible morphological forms $y_t^{l_i}$ (i.e., $P(y_t^{l_i} | l_i)$). Specifically, we use \mathbf{h}_t to choose an inflected form from its possible inflected forms (e.g., Singular or Plural in English):

$$P(y_t^{l_i} | l_i) = \text{softmax}(\mathbf{W}_{l_i} \mathbf{h}_t) \quad (7)$$

Here $\mathbf{W}_{l_i} \in \mathbb{R}^{s_i \times d}$ where s_i is the number of inflected forms of label l_i (in most cases 2 for English, singular and plural). Finally, the **ECOL** model concatenates the above refined probabilities as follows:

$$P(y_t^v) = v_t \quad (8)$$

$$P(y_t^{l_i}) = c_{t,i} \cdot P(y_t^{l_i} | l_i) \quad (9)$$

$$P(y_t) = P(y_t^v) \odot P(y_t^{l_1}) \odot \dots \odot P(y_t^{l_{k^f}}) \quad (10)$$

where \odot represents concatenation. Some novel object labels are included in the GPT2 vocabulary.

However, these words are not present in the training captions and thus the model always assigns them very low probabilities in $P(y_t^v)$. The only way novel object labels can appear in captions is through copy operations.

Model Application Scope In this paper, we focus on the Novel Object Captioning task. However, in general, our copy mechanism is capable of copying any type of information. The Abstract Label approach is general to zero shot learning problems where novel items share characteristics with training instances. The Morphological Selector is also applicable to linguistic copy mechanisms in other contexts such as Commonsense Reasoning (Lin et al., 2020) where copied terms may require linguistic alignment with the generated text.

3.2 Copying More Object Labels

In this paper, we encourage the copying of object labels by using a suitable reward function in the Self-Critical Sequence Training (SCST) framework, which has proven effective for image captioning tasks. Compared with injecting additional loss terms together with the standard XE loss, using the SCST framework allows us to design arbitrary encouragement signals based on the inference output. It minimizes the negative expected reward score:

$$L_R(\theta) = -\mathbf{E}_{y_{1:T} \sim p_\theta} [r(y_{1:T})] \quad (11)$$

where r is the reward function and p_θ represents the models outputs. In this paper, following Cornia et al. (2020), we first pre-train the **ECOL** model with the CE loss, then switch to fine-tune the **ECOL** model with the above SCST loss.

Inference Bias Baseline: We add an Inference Bias (IB) $b \in \mathbb{R}^+$ to increase $P(y_t^{l_i})$ at inference time. Eq. 9 is changed to:

$$P(y_t^{l_i}) = b \cdot c_{t,i} \cdot P(y_t^{l_i} | l_i) \quad (12)$$

and remaining probabilities normalised accordingly. IB is functionally equivalent to adjusting the threshold for the copy decision during inference. Surprisingly, this simple inference trick provides a strong baseline (see Table 3). This shows that after the CE training, many correct copy operations are assigned with low probabilities, compared to the fixed vocabulary items. However, we believe that it is better to train the model to increase the probabilities of these copy operations than adding ad hoc adjustments at inference time.

Can Standard SCST Encourage Copying? Rennie et al. (2017) shows that SCST with the CIDEr reward fine-tuning leads to noticeable caption quality improvement for standard image captioning tasks (i.e., improvement in various automatic metrics). Previous work (Cornia et al., 2020; Anderson et al., 2018b) use CIDEr as the standard reward function in their SCST optimization. This shows suggests that the problem of overfitting of SCST training with CIDEr reward is minimal. Intuitively, the CIDEr reward is positively correlated with the number of salient object label mentions and should encourage the model to copy salient novel object labels. However, CIDEr equally rewards both generation of object labels present in training data via the vocabulary $P(y_t^v)$ and via copy operations $P(y_t^l)$. Novel objects labels however can only be generated by copy operations (see Sec. 3.1), thus the CIDEr reward function does not sufficiently encourage these operations. We propose two orthogonal modifications to the standard SCST algorithm to address this issue:

Novel Encouragement Reward: We propose combining the standard CIDEr-D reward with a reward function that gives captions with words copied from object labels an extra bonus, which we intend to encourage copy operations. One straightforward way to implement this idea is to provide a constant bonus to each triggered copy operation:

$$R_a(X) = CIDEr(X) + a * C \quad (13)$$

where X is a generated caption, C is the number of copy actions in the caption X and $a \in R^+$ is a fixed hyper-parameter. We refer this as *additive bias*. Optimizing with the additive bias, the captioning model only needs to trigger the copy operation for arbitrary objects at arbitrary generation steps. That is, the model may encourage copying object labels at the expense of caption quality (i.e., high CIDEr-D scores). Therefore, we propose a *proportional bias* that assigns different rewards to the copy operations in different images by making a connection between the copy bonus and the generated captions CIDEr-D score:

$$R_p(X) = CIDEr(X) * (1.0 + p * C) \quad (14)$$

where $p \in R^+$ is a fixed hyper-parameter. Although R_a can effectively encourage the model to copy objects, it may introduce noisy object mentions. R_p penalizes those noisy object mentions via the low caption CIDEr score.

Visual Object Aligned (VOA) Images: *VOA Images* refers to the set of training images where the reference captions contain at least one word from retained object labels. During SCST training, images that contain no object label words (i.e., non-VOA images) will not utilise copy operations, thus these images encourage the model NOT to copy. VOA images account for approximately 70% of the full COCO-2017 training images set. Although restricting training to VOA images can be done with arbitrary models, this may hurt the diversity of generated captions. Experiments in Table 3 confirm that restricting to VOA images only improves performance when used with SCST training.

Hyper-Parameters For Copy Encouragement: The above approaches introduce two additional parameters: a and p . In our experiments, a and p range over 0.2, 0.3 and 0.4; we found that 0.3 works the best for both reward functions. Combined with restricting SCST training to VOA images, R_p works better than R_a and sets a new SOTA for novel object image captioning.

4 Experiments

We conduct experiments on the *nocaps* (Agrawal et al., 2019) and the *held-out* COCO (Hendricks et al., 2016) Benchmark. We set the layer and embedding size to $d = 768$ and use Adam optimisation (Kingma and Ba, 2014). We train our models 15 epochs with batch size 100 for CE loss and 15 epochs with batch size 10 for SCST loss.

4.1 Evaluation Metrics

We use CIDEr (Vedantam et al., 2015), SPICE (Anderson et al., 2016a) and METEOR (Banerjee and Lavie, 2005) to evaluate the caption quality. CIDEr measures the similarity between the reference captions and generated outputs using tf-idf weighted n-gram overlap. SPICE is based on the scene graphs matching between the reference captions and generated outputs. METEOR focuses on the alignment between the words in reference captions and generated outputs, with an aim of 1:1 correspondence. To measure the effectiveness of our copy encouragement approach, we report object F1 (Anderson et al., 2017) in the *held-out* COCO Benchmark. As the *nocaps* benchmark does not release its ground-truth captions, we instead report averaged number of mentioned objects (Ave. O) and CIDEr score for dummy captions that only contain copied object words (Object CIDEr, OC., details see Appendix).

Method	<i>in-domain</i>		<i>near-domain</i>		<i>out-of-domain</i>		Meteor	Overall	
	CIDEr	SPICE	CIDEr	SPICE	CIDEr	SPICE		CIDEr	SPICE
Up-Down + BS	73.7	11.6	57.2	10.3	30.4	8.1	22.9	54.5	10.1
Up-Down + ELMo + CBS	76.0	11.8	74.2	11.5	66.7	9.7	24.4	73.1	11.2
NBT + BS	62.8	10.3	51.9	9.4	48.9	8.4	21.8	54.3	9.4
NBT + CBS	61.9	10.4	57.3	9.6	61.8	8.6	21.6	59.9	9.5
<i>OSCAR_L</i> + CBS + SCST	-	-	-	-	-	-	-	80.9	11.3
ECOL + IB (Ours)	81.7	12.9	77.2	12.1	67.0	10.3	25.6	76.0	11.9
ECOL-R (Ours)	87.3	12.8	84.0	12.5	75.4	10.7	25.7	82.9	12.2

Table 1: Comparison of the **ECOL-R** model with other state-of-the-art systems on the *nocaps* Test Split. The **ECOL-R** model sets new state of the art and improves previous work by 2.0 CIDEr and 0.9 SPICE. The performance of the **Up-Down** and **NBT** models are from Agrawal et al. (2019). The *OSCAR_L* model is from Li et al. (2020).

Method	<i>out-of-domain</i>				<i>in-domain</i>		
	Meteor	CIDEr	SPICE	Object F1	Meteor	CIDEr	SPICE
LSTM-P (Li et al., 2019)	23.4	88.3	16.6	60.9	-	-	-
Base + CBS (Anderson et al., 2017)	23.3	77.0	15.9	54.0	24.9	88.0	18.4
NBT + CBS (Lu et al., 2018a)	24.1	86.0	17.4	70.5	25.0	92.1	18.0
ECOL + IB (Ours)	25.6	95.5	18.8	58.2	27.0	108.3	20.4
ECOL-R (Ours)	25.7	99.2	19.3	66.3	26.8	113.3	20.4
ECOL-R + CBS (Ours)	25.7	99.1	19.1	71.8	26.8	112.6	20.8
PS3 (Anderson et al., 2018a) [§]	25.4	94.5	17.9	63.0	25.9	101.1	19.0
FDM-net (Cao et al., 2020) [§]	25.9	84.8	19.4	64.7	27.2	109.7	20.2
FDM-net + CBS (Cao et al., 2020) [§]	25.6	85.3	19.6	85.7	26.2	105.5	19.7

Table 2: Comparison of the **ECOL-R** model with previous work on the test split of the *held-out* COCO Benchmark. Object F1 measures the presence of particular novel objects in the captions. [§] not comparable as they use additional information related to novel objects (i.e., images containing novel objects and scene graphs).

4.2 Comparison with the State-of-the-art

We compare our models **ECOL** + IB and **ECOL-R** with other state-of-the-art systems in Tables 1 and 2.

On the *nocaps* benchmark (Table 1), our models outperform previous work, including the recently proposed *OSCAR_L* + CBS + SCST model (Li et al., 2020), which is fine-tuned from the BERT-LARGE model (Devlin et al., 2019), by 2.0 CIDEr, 0.9 SPICE and set a new state of the art. Compared with the *OSCAR_L* model, our models use far fewer model parameters (340M vs. 60M) and outperforms *OSCAR_L* on both CIDEr and SPICE metrics. We train our model for about 10 hours for CE Loss and 24 hours for SCST Loss using a single Nvidia P100 GPU. As a comparison, the *OSCAR_L* model which is fine-tuned from BERT-LARGE uses 60 - 90 hours for training CE Loss

and 60 - 90 hours for training SCST Loss.² This shows that simply deploying a BERT-based language model is not sufficient for the Novel Object Captioning task.

On the *held-out* COCO benchmark (Table 2), the **ECOL-R** model produces more novel objects (+ 13.3 Object F1) and higher quality captions (+ 3.9 CIDEr on the *out-of-domain* split) than the **ECOL** model with run-time Inference Bias. Compared with previous work, the **ECOL-R** model achieves 10.9 CIDEr and 1.9 SPICE higher in the *out-of-domain* split, 21.2 CIDEr and 2.8 SPICE higher in the *in-domain* split with the highest object F1. This shows that our copy encouragement approach successfully trains our model to correctly copy more novel objects and to produce high-quality captions. Compared with PS3 (Anderson et al., 2018a)

²According to the authors' comments on their official model code repo <https://github.com/microsoft/Oscar/issues/6>

and FDM-net model (Cao et al., 2020) which are trained on extra images containing novel objects and scene graphs, our models still outperform the PS3 model and 13.9 CIDEr higher than the FDM-net. We set a new state of the art in this benchmark without additional novel objects information.

4.3 Ablation Study

Table 3 presents ablation results for various **ECOL-R** components, including our copy encouragement approach. Table 4 shows that our encouragement of copying in the **ECOL-R** model does not benefit from additional Inference Bias. Table 5 shows the effect of Abstract Labels and the Morphological Selector in the **ECOL-R** model. Finally, Table 6 confirms the **ECOL-R** model’s generalization ability for in-domain COCO images.

ECOL-R Components: The **ECOL** model produces better captions using the frozen GPT2 parameters (row 1 vs. 2). Our copy mechanism (C) helps the model to explicitly integrate novel objects, substantially improving the *out-of-domain* split by 15.3 CIDEr and 0.3 SPICE (row 2 vs. 3). The Inference Bias (IB) introduces noticeable performance improvement: 8.4 CIDEr and 0.3 SPICE (row 3 vs. 4) in models that do not use our reinforcement learning approach. The **ECOL** model trained with the standard SCST reward function obtains an overall 8.1 CIDEr improvement, but most of the improvement is from the *in-domain* and *near-domain* splits (row 8 vs. 6). Compared with the **ECOL + IB** model, the **ECOL** model trained with standard SCST algorithm is 8.1 CIDEr lower in the *out-of-domain* split (row 5 vs. 4). As discussed in Sec. 3.2, standard SCST cannot provide sufficient copy encouragement as object words can be generated from either pathways (fixed vocabulary or copy). Optimizing either the R_a or R_p reward functions improves the **ECOL + CIDEr** model by 7.0 CIDEr and 7.8 CIDEr respectively (row 7 and 5). R_a achieves 3.7 CIDEr higher than R_p in the *out-of-domain* split. Interestingly, after restricting the model training to VOA images, R_p achieves 7.8 CIDEr improvement in the *out-of-domain* split (row 8 vs. 7), outperforming the **ECOL + R_a** w/ VOA model by 1.4 CIDEr (row 10 vs. 8).

Effectiveness of Copy Encouragement: We directly measure the copy quantity by counting the number of copied object labels and Object CIDEr. Row 5 and 3 confirm that the standard SCST algorithm has little impact on the copy quantity (only +

0.1 object per image and + 1.4 Object CIDEr). Inference Bias (IB), R_a and R_p rewards substantially improve the quantity of copied objects (row 4, 6, 7 vs. 3). Among these three components, the models trained with R_a and R_p work better than the IB baseline (row 6, 7 vs. 4). The model trained with the R_a reward copies more objects than the R_p reward, especially training with all training images. This is because the R_a reward assigns constant positive reward for all copied objects. However, such a naive reward appears to encourage noisy copying operations (i.e., copying non-salient objects). As a result, the **ECOL + R_a** model performs worse than the **ECOL + R_p** model in terms of caption quality (row 7 vs. 6). After restricting training with VOA images, the models trained with R_a and R_p copy similar amount of objects, but the model with R_a produce better captions than the one with R_p , especially in the *out-of-domain* split (row 10 vs. 8). The R_p reward maintains a good balance between copying more objects and high caption quality.

Are The VOA images Always Useful? Restricting training to the VOA images can be done with any captioning models. However, this does not necessarily encourage copy operations and improve the output caption quality. When we restrict training to VOA images, the **ECOL-R** model performs consistently worse in all three splits compared to our proposed training scheme (row 9 vs. 10). The only difference is that the **ECOL** model is not trained with diverse images during the cross-entropy stage. That is, restricting to VOA images is only suitable for fine-tuning in the SCST stage.

Sufficient Encouragement For Copy: Here we investigate whether our **ECOL-R** model mentions a sufficient number of salient objects. We apply an increasing amount of inference bias to the **ECOL**, **ECOL + CIDEr** and **ECOL-R** models in Table 4. We note that only **ECOL-R** model is negatively impacted (measured by CIDEr score) by different Inference Bias values. This shows that the **ECOL-R** model does not benefit from further copy encouragement.

The Effect of Abstract Labels and M Selector Table 5 shows the effect of Abstract Labels (AL) and the M Selector (M) in the **ECOL + IB** and **ECOL-R** models. Removing AL and M from the **ECOL + IB** model drops 2.3 CIDEr, 0.6 SPICE and 4.6 CIDEr, 0.8 SPICE respectively. AL and M have a large positive impact on the SPICE

Method	<i>in-domain</i>		<i>near-domain</i>		<i>out-of-domain</i>		Overall			
	CIDEr	SPICE	CIDEr	SPICE	CIDEr	SPICE	CIDEr	SPICE	OC.	Ave. O
<i>Ablation of ECOL Components</i>										
(1) R + P + AL	80.3	12.7	60.3	11.4	34.8	9.2	58.0	11.2	-	-
(2) GPT2 + R + P + AL	78.6	12.5	64.2	11.6	39.4	9.3	61.3	11.3	-	-
(3) C + GPT2 + R + P + AL (ECOL)	80.7	12.7	68.6	11.8	61.0	10.2	68.8	11.7	10.3	0.5
<i>Ablation of Copy Encouragement</i>										
(4) ECOL + IB	85.3	12.8	76.6	12.2	73.7	10.8	77.2	12.0	17.5	1.1
(5) ECOL + CIDEr	89.7	12.3	75.9	11.8	63.5	10.1	75.3	11.5	11.7	0.6
(6) ECOL + R_a	89.0	12.6	82.2	12.5	78.2	11.2	82.3	12.3	19.7	1.5
(7) ECOL + R_p	92.3	12.9	83.5	12.5	74.9	11.0	83.0	12.3	17.7	1.3
(8) ECOL + R_a w/ VOA	89.3	12.6	83.6	12.5	81.7	11.1	84.0	12.2	20.7	1.6
(9) ECOL + R_p w/ VOA (all training)	88.0	12.6	83.3	12.6	79.4	11.2	83.2	12.3	20.6	1.5
(10) ECOL + R_p w/ VOA (ECOL-R)	90.4	12.6	84.5	12.4	83.1	11.0	85.1	12.2	20.4	1.5

Table 3: Ablation study of our model. *OC.* for Object CIDEr; *Ave. O* for Averaged Number of Mentioned Object in each image; *C* for Copy Mechanism; *R* for ROI; *P* for Position; *AL* for Abstract Label; R_a and R_p are the SCST reward function; *VOA*: Visual Object Aligned Images; *VOA (all training)*: all training using VOA images.


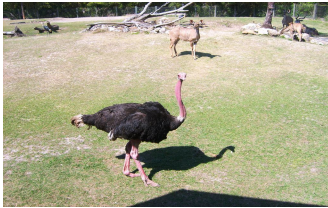

	<i>curtain</i>	<i>ostrich, deer</i>	<i>door, house</i>
ECOL-R	 A bathroom with a shower <u>curtain</u> and a toilet. ✓	 <u>An ostrich</u> and a deer standing in a field. ✓	 A red door of a red house with a red phone. ✗
ECOL + IB	A white bath tub sitting next to a white toilet. ✗	Two ostriches and a deer in a grassy field. ✗	A red telephone booth sitting next to a brick wall. ✓
GT	The bathtub is white and has a white shower <u>curtain</u> .	<u>An ostrich</u> standing in grass with a few deer in the background.	A red phone booth is standing against a brick wall.

Figure 4: Three Examples generated by the **ECOL-R** and **ECOL + IB** model on the *nocaps* Val Set. First Case: The **ECOL-R** model accurately mentions the novel object *curtain* in the caption. Second Case: Both models talk about *ostrich* in their generated caption, but the **ECOL + IB** model uses the wrong modifier. Third Case: when detected object labels are too general, the **ECOL-R** model may produce inaccurate captions.

IB	1 (e^0)	2.72 (e^1)	7.39 (e^2)	20.01 (e^3)
ECOL	68.8	73.6	76.0	77.2
ECOL + CIDEr	75.3	77.6	79.5	81.1
ECOL-R	85.1	85.0	84.9	84.3

Table 4: The CIDEr Score in *nocaps* Validation Set with different Inference Bias.

score. As SPICE is sensitive to long-range object word relationships, such as attributes and predicate words, (Anderson et al., 2016a) Abstract Labels and the M Selector improve the semantic coherence and fluency of the captions. The performance gap in the **ECOL-R** model becomes smaller. Our copy encouragement approach contributes to the generation coherency and fluency.

Method	<i>out-of-domain</i>		Overall	
	CIDEr	SPICE	CIDEr	SPICE
ECOL + IB	73.7	10.8	77.2	12.0
ECOL + IB w/o AL	72.7	10.0	74.9	11.4
ECOL + IB w/o M	71.2	10.5	72.6	11.2
ECOL-R	83.1	11.0	85.1	12.2
ECOL-R w/o AL	80.8	10.9	84.4	12.1
ECOL-R w/o M	78.4	10.7	82.8	11.9

Table 5: The contribution of Abstract Label (AL) and Morphological Selector (M).

Generalization For In-Domain COCO: To further show the generalization of our model for In-Domain images (i.e., without novel objects), we run the **ECOL + IB** and **ECOL-R** models on the COCO 2017 Validation Set and compare with another two

Method	Meteor	CIDEr	SPICE
NBT + CBS	25.1	80.2	15.8
Up-Down + CBS	25.7	95.4	18.2
ECOL + IB	27.1	106.1	20.6
ECOL-R	27.1	114.2	20.9

Table 6: The performance of the best four *nocaps* models on the COCO 2017 Validation Set.

novel object captioning models (**NBT** + CBS and **Up-Down** + ELMo + CBS) reported in Agrawal et al. (2019) in Table 6. Both of our models outperforms the **Up-Down** and **NBT** model by a large margin. Our models produce high-quality captions for images with novel objects as well as known objects.

4.4 Qualitative analysis on *nocaps*

Qualitative analysis on the *nocaps* validation set reveals that the **ECOL-R** model mentions the salient object in the input image (first example in Figure 4), is able to generate more accurate descriptions of novel objects (second example in Figure 4), however may generate inaccurate captions due to the non-informative detected object labels (third example in Figure 4). In summary, the **ECOL-R** model is better at incorporating detected image objects into generated captions than the **ECOL** + IB model.

5 Conclusion and Future work

This paper proposes the **ECOL-R** model that includes a training scheme to encourage copying novel object labels using Reinforced Learning. Our experiments show that the **ECOL-R** model successfully integrates novel object information and achieves state-of-the-art performance on two Novel Object Caption Benchmarks. In the future, we plan to extend our SCST reward function to other metrics such as SPICE (Anderson et al., 2016b) and BertScore (Zhang et al., 2020).

Acknowledgments

We thank anonymous reviewers for their insightful suggestions to improve this paper. This research was supported by a Google award through the Natural Language Understanding Focused Program, by a MQ Research Excellence Scholarship and a CSIRO’s DATA61 Top-up Scholarship, and under the Australian Research Councils Discovery Projects funding scheme (project number

DP160102156).

References

- Harsh Agrawal, Karan Desai, Yufei Wang, Xinlei Chen, Rishabh Jain, Mark Johnson, Dhruv Batra, Devi Parikh, Stefan Lee, and Peter Anderson. 2019. *nocaps: novel object captioning at scale*. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*.
- Peter Anderson, Basura Fernando, Mark Johnson, and Stephen Gould. 2016a. *Spice: Semantic propositional image caption evaluation*. In *ECCV*.
- Peter Anderson, Basura Fernando, Mark Johnson, and Stephen Gould. 2016b. *Spice: Semantic propositional image caption evaluation*. In *European Conference on Computer Vision*, pages 382–398. Springer.
- Peter Anderson, Basura Fernando, Mark Johnson, and Stephen Gould. 2017. *Guided open vocabulary image captioning with constrained beam search*. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 936–945, Copenhagen, Denmark. Association for Computational Linguistics.
- Peter Anderson, Stephen Gould, and Mark Johnson. 2018a. *Partially-supervised image captioning*. In *Advances in Neural Information Processing Systems*, pages 1875–1886.
- Peter Anderson, Xiaodong He, Chris Buehler, Damien Teney, Mark Johnson, Stephen Gould, and Lei Zhang. 2018b. *Bottom-up and top-down attention for image captioning and visual question answering*. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Satanjeev Banerjee and Alon Lavie. 2005. *Meteor: An automatic metric for mt evaluation with improved correlation with human judgments*. In *Proceedings of the acl workshop on intrinsic and extrinsic evaluation measures for machine translation and/or summarization*, pages 65–72.
- Tingjia Cao, Ke Han, Xiaomei Wang, Lin Ma, Yanwei Fu, Yu-Gang Jiang, and Xiangyang Xue. 2020. *Feature deformation meta-networks in image captioning of novel objects*. In *Proceedings of the AAAI Conference on Artificial Intelligence*.
- Massimiliano Ciaramita and Mark Johnson. 2003. *Supersense tagging of unknown nouns in WordNet*. In *Proceedings of the 2003 Conference on Empirical Methods in Natural Language Processing*, pages 168–175.
- Marcella Cornia, Matteo Stefanini, Lorenzo Baraldi, and Rita Cucchiara. 2020. *Meshed-Memory Transformer for Image Captioning*. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*.

- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. [BERT: Pre-training of deep bidirectional transformers for language understanding](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4171–4186, Minneapolis, Minnesota. Association for Computational Linguistics.
- Lisa Anne Hendricks, Subhashini Venugopalan, Marcus Rohrbach, Raymond J. Mooney, Kate Saenko, and Trevor Darrell. 2016. [Deep compositional captioning: Describing novel object categories without paired training data](#). In *2016 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2016, Las Vegas, NV, USA, June 27-30, 2016*, pages 1–10.
- Simao Herdade, Armin Kappeler, Kofi Boakye, and Joao Soares. 2019. [Image captioning: Transforming objects into words](#). In H. Wallach, H. Larochelle, A. Beygelzimer, F. dAlché Buc, E. Fox, and R. Garnett, editors, *Advances in Neural Information Processing Systems 32*, pages 11137–11147. Curran Associates, Inc.
- Ronghang Hu, Amanpreet Singh, Trevor Darrell, and Marcus Rohrbach. 2020. Iterative answer prediction with pointer-augmented multimodal transformers for textvqa. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Diederik P Kingma and Jimmy Ba. 2014. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*.
- Alina Kuznetsova, Hassan Rom, Neil Alldrin, Jasper Uijlings, Ivan Krasin, Jordi Pont-Tuset, Shahab Kamali, Stefan Popov, Matteo Mallocci, Tom Duerig, et al. 2018. The open images dataset v4: Unified image classification, object detection, and visual relationship detection at scale. *arXiv preprint arXiv:1811.00982*.
- Xiujun Li, Xi Yin, Chunyuan Li, Xiaowei Hu, Pengchuan Zhang, Lei Zhang, Lijuan Wang, Houdong Hu, Li Dong, Furu Wei, Yejin Choi, and Jianfeng Gao. 2020. Oscar: Object-semantics aligned pre-training for vision-language tasks. *ECCV*.
- Yehao Li, Ting Yao, Yingwei Pan, Hongyang Chao, and Tao Mei. 2019. Pointing novel objects in image captioning. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Bill Yuchen Lin, Wangchunshu Zhou, Ming Shen, Pei Zhou, Chandra Bhagavatula, Yejin Choi, and Xiang Ren. 2020. [CommonGen: A constrained text generation challenge for generative commonsense reasoning](#). In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 1823–1840, Online. Association for Computational Linguistics.
- Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C. Lawrence Zitnick. 2014. Microsoft coco: Common objects in context. In *Computer Vision – ECCV 2014*, pages 740–755, Cham. Springer International Publishing.
- Jiasen Lu, Jianwei Yang, Dhruv Batra, and Devi Parikh. 2018a. Neural baby talk. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 7219–7228.
- Yujie Lu, Boyi Ni, Qijin Ji, Kotaro Sakamoto, Hideyuki Shibuki, and Tatsunori Mori. 2018b. [Deep learning paradigm with transformed monolingual word embeddings for multilingual sentiment analysis](#). In *Proceedings of the 32nd Pacific Asia Conference on Language, Information and Computation*, Hong Kong. Association for Computational Linguistics.
- Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, and Ilya Sutskever. 2019. Language models are unsupervised multitask learners. *OpenAI Blog*, 1(8):9.
- Marc’Aurelio Ranzato, Sumit Chopra, Michael Auli, and Wojciech Zaremba. 2015. Sequence level training with recurrent neural networks. *arXiv preprint arXiv:1511.06732*.
- Steven J. Rennie, Etienne Marcheret, Youssef Mroueh, Jerret Ross, and Vaibhava Goel. 2017. Self-critical sequence training for image captioning. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. [Attention is all you need](#). In I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett, editors, *Advances in Neural Information Processing Systems 30*, pages 5998–6008. Curran Associates, Inc.
- Ramakrishna Vedantam, C. Lawrence Zitnick, and Devi Parikh. 2015. Cider: Consensus-based image description evaluation. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Oriol Vinyals, Meire Fortunato, and Navdeep Jaitly. 2015. [Pointer networks](#). In C. Cortes, N. D. Lawrence, D. D. Lee, M. Sugiyama, and R. Garnett, editors, *Advances in Neural Information Processing Systems 28*, pages 2692–2700. Curran Associates, Inc.
- Weixuan Wang, Zhihong Chen, and Haifeng Hu. 2019. Hierarchical attention network for image captioning. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, pages 8957–8964.
- Yu Wu, Linchao Zhu, Lu Jiang, and Yi Yang. 2018. [Decoupled novel object captioner](#). In *2018 ACM Multimedia Conference on Multimedia Conference, MM 2018, Seoul, Republic of Korea, October 22-26, 2018*, pages 1029–1037.

Deshraj Yadav, Rishabh Jain, Harsh Agrawal, Prithvijit Chattopadhyay, Taranjeet Singh, Akash Jain, Shiv Baran Singh, Stefan Lee, and Dhruv Batra. 2019. Evalai: Towards better evaluation systems for ai agents. *arXiv preprint arXiv:1902.03570*.

Ting Yao, Yingwei Pan, Yehao Li, and Tao Mei. 2017. Incorporating copying mechanism in image captioning for learning novel objects. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 6580–6588.

Tianyi Zhang, Varsha Kishore, Felix Wu, Kilian Q. Weinberger, and Yoav Artzi. 2020. **Bertscore: Evaluating text generation with bert**. In *International Conference on Learning Representations*.

Appendices

A Model Details

The hyper-parameters of the **ECOL-R** model is shown in Table 7. This architecture is basically from (Cornia et al., 2020). We only change the hidden size of the model to 768 to fit the size of GPT2 (the smallest version). Our model has total 60.8×10^6 parameters and 43.0×10^6 trainable parameters. This scale is slightly smaller than the Transformer Base model (65.8×10^6) (Vaswani et al., 2017).

We optimise with Adam($\alpha=0.9$, $\beta=0.98$, $\epsilon=1e-9$) (Kingma and Ba, 2014) and clip gradients to 0.1 for both Benchmarks. In Cross-entropy training, we vary the learning rate over the course of training using the heuristic:

$$lrate = d_{\text{model}}^{-0.5} \cdot \min(S^{-0.5}, S \cdot W^{-1.5}) \quad (15)$$

where S is the step number and W is the number of warm-up steps. We set W to 20000 steps for the *nocaps* Benchmark and 10000 steps for the *held-out* COCO Benchmark. The number of warm-up steps has some impact on both benchmark. We tried 20,000 and 10,000 for both Benchmarks. For SCST training, we set the initial learning rate $1e^{-6}$ and reduce it by half if the reward metric (Validation set CIDEr) does not improve for 3 evaluations. We conduct evaluation every 3000 steps.

We use Pytorch 1.4.0 to implement our model. The Cross-Entropy Training takes about 8 hours and the SCST optimization takes about 15 hours in a single NVIDIA Tesla P100 GPU.

Our source code is submitted in the *Software*. We setup an anonymous Google Drive to host large file³.

Name	Value	Name	Value
N_e	3	Dropout	0.1
N_d	3	Max Cap. Len.	20
Model Size	768	beam Size	5
FFN Size	2048	Att. Head Num	8

Table 7: Hyper-parameters of the **ECOL-R** model.

A.1 Input Object Detector

We follow the processing of input objects in Agrawal et al. (2019). We observed that some

³<https://drive.google.com/drive/folders/1EToBXQ8WAWxn5uCd38HtFRYmchnBBbMo?usp=sharing>

object categories are frequently mentioned in the training captions and that they often have variable, context-sensitive verbalisation (e.g., a person might be described as a sports player, a student, etc., depending on the context). For those objects, vocabulary based word generation often did a better job at selecting the correct verbalisation due to their frequency in training captions. On the other hand, novel objects typically have lower-frequencies and a fixed, single verbalisation. For example, elephants are usually only referred to with the word *elephant*. For this reason, we remove objects with high-frequency in training captions from the output of the task specific object detector, leaving their corresponding words to be generated via vocabulary softmax. We also remove the more abstract objects (higher in the object class hierarchy) when object regions overlap. Finally, we keep only one detected object for each label (the one with highest confidence score).

We provide the downloadable link of filtered results in Sec B. We use exactly the same Visual Genome objects as described in Anderson et al. (2018b). The Visual Genome object detector (Anderson et al., 2018b) can produce ROI vectors for arbitrary bounding boxes, hence we use it also to produce ROI vectors for objects from the task specific detector.

A.2 ECOL-R Inference Details

We use Beam Search with beam size 5 to decode the captions. We first do length normalization for the overall score of each decoded caption. We also penalize captions when they generate repeated bi-grams. Once the repetitions are found, the log-probability for that particular word is divided by a penalty value e^2 . All image objects are only allowed to be copied once. During the SCST optimization, we mask out words from the vocabulary that can be generated via copy operations to encourage the model to copy. All the above constraints are applied to all of our models in the ablation study.

A.3 Object CIDEr Details

Object CIDEr score for dummy captions that only contain copied object words. This shows the correctness of our copy mechanism. High Object CIDEr score means many of the copied object labels are also mentioned in the ground-truth captions. We use this metrics because the *nocaps* benchmark does not release its ground-truth captions and only provide online evaluation APIs.

B Dataset Details

For the *nocaps* Benchmark, we train with the COCO-2017 dataset, which is available at <http://images.cocodataset.org/zips/train2017.zip>. The *nocaps* Validation and Test datasets are available from <https://nocaps.org/download>. The Visual Genome image object detection files can be found in <https://github.com/nocaps-org/updown-baseline>. For the *held-out* COCO Benchmark, the training and evaluation data can be found in <https://github.com/LisaAnne/DCC>. The Visual Genome image object detector is used for both benchmarks because COCO-2017 and COCO-2014 share the same set of images. The anonymous Google Drive includes the above data and the sets of task specific objects detected for the above two benchmarks.

B.1 Duplicated Caption Removal

We find some images in COCO share exactly the same reference captions. We find it beneficial to remove those duplicates. We simply iterate over all reference captions and remove any captions if they have already been found previously. This removes 25463 captions from the training data of the *nocaps* Benchmark and 7059 captions from the training data of the *held-out* COCO Benchmark.

B.2 VOA (visual object aligned) Images

VOA (visual object aligned) images/reference caption pairs are those that mention at least one detected task specific image object label (or their linguistic variant). Non-VOA image/caption pairs are removed from our SCST training process. We provide the reduced set of reference captions in the anonymous Google Drive (*ddc_captions/ddc_train_VOA.json* and *nocaps_captions/nocaps_train_VOA.json*).

B.3 Data Statistics for Benchmarks

Table 9 and Table 10 show the number of images and annotated reference captions of the *nocaps* and *held-out* COCO Benchmark, respectively. On average, each image has five reference captions. The COCO Train in the *nocaps* Benchmark is larger than the *held-out* COCO Benchmark.

C Evaluation

The *nocaps* Benchmark hosts its evaluation sever at <https://evalai.cloudcv.org/web/>

Method	<i>out-of-domain</i>				<i>in-domain</i>		
	Meteor	CIDEr	SPICE	Object F1	Meteor	CIDEr	SPICE
ECOL + IB (Ours)	25.1	93.0	18.7	58.4	26.9	107.1	20.3
ECOL-R (Ours)	25.6	99.2	19.5	64.5	26.7	113.0	20.3
ECOL-R + CBS (Ours)	25.6	98.4	19.3	70.9	26.7	112.3	20.2

Table 8: The performance of **ECOL + IB**, **ECOL-R** and **ECOL-R + CBS** model on *held-out* COCO Benchmark Validation Set.

	COCO Train	COCO Train VOA	COCO Val	<i>nocaps</i> Val	<i>nocaps</i> Test
#Image	118,287	82,771	5,000	4,500	10,600
#Caption	591,753	299,502	25,014	-	-

Table 9: Data Statistics for the *nocaps* Benchmark. The full set of annotations in *nocaps* Val and Test is not available. One can only access some of them via <https://nocaps.org/explore>.

	Train	Train VOA	Val in-domain	Val out-domain	Test in-domain	Test out-domain
#Image	70,194	55,799	17,234	3,018	17,288	3,024
#Caption	351,134	197,061	86,230	1,5105	86,188	15,131

Table 10: Data Statistics for the *held-out* COCO Benchmark.

[challenges/challenge-page/355/overview](https://nocaps.org/challenges/challenge-page/355/overview).

The detailed setup instruction of the local submission environment to Evai (Yadav et al., 2019) is available at <https://github.com/nocaps-org/updown-baseline>.

The *held-out* COCO Benchmark uses the evaluation tool from <https://github.com/ruotianluo/coco-caption/tree/ea20010419a955fed9882f9dcc53f2dc1ac65092>.

We provide an on-the-shelf version of this tool in the anonymous Google Drive (in *tools*).

C.1 *held-out* COCO Benchmark Validation Performance

We only show the test performance on the *held-out* COCO Benchmark in our main paper. Here, we show the performance of our model performance on the validation Set in Table 8. The models achieve similar level of performance on the Validation Set.