

Reinforcing an Image Caption Generator Using Off-Line Human Feedback

Paul Hongsuck Seo,^{1,3*} Piyush Sharma,² Tomer Levinboim,² Bohyung Han,³ Radu Soricut²

¹Computer Vision Lab., POSTECH, Korea

²Google Research, USA

³Computer Vision Lab., ECE, & ASRI, Seoul National University, Korea

hsseo@postech.ac.kr, {piyushsharma, tomerl, rsoricut}@google.com, bhhan@snu.ac.kr

Abstract

Human ratings are currently the most accurate way to assess the quality of an image captioning model, yet most often the only used outcome of an expensive human rating evaluation is a few overall statistics over the evaluation dataset. In this paper, we show that the signal from *instance-level* human caption ratings can be leveraged to improve captioning models, even when the amount of caption ratings is several orders of magnitude less than the caption training data. We employ a policy gradient method to maximize the human ratings as rewards in an off-policy reinforcement learning setting, where policy gradients are estimated by samples from a distribution that focuses on the captions in a caption ratings dataset. Our empirical evidence indicates that the proposed method learns to generalize the human raters' judgments to a previously unseen set of images, as judged by a different set of human judges, and additionally on a different, multi-dimensional side-by-side human evaluation procedure.

1 Introduction

Image captioning is the task of automatically generating fluent natural language descriptions for an input image. However, measuring the quality of generated captions in an automatic manner is a challenging and yet-unsolved task; therefore, human evaluations are often required to assess the complex semantic relationships between a visual scene and a generated caption (Sharma et al. 2018; Cui et al. 2018; Zhao et al. 2019). As a result, there is a mismatch between the training objective of the captioning models and their final evaluation criteria. The most simple and frequently-used training objective is maximum likelihood estimation (MLE) (Vinyals et al. 2015; Mun, Cho, and Han 2017; Lu et al. 2017; 2018; Changpinyo et al. 2019), while other approaches make use of handcrafted evaluation metrics, such as CIDEr (Vedantam, Zitnick, and Parikh 2015), to optimize model parameters using reinforcement learning (RL) (Rennie et al. 2017; Ding and Soricut 2017; Anderson et al. 2018; Qin et al. 2019). However, these surrogate objectives capture only limited aspects of caption quality, and often fail to

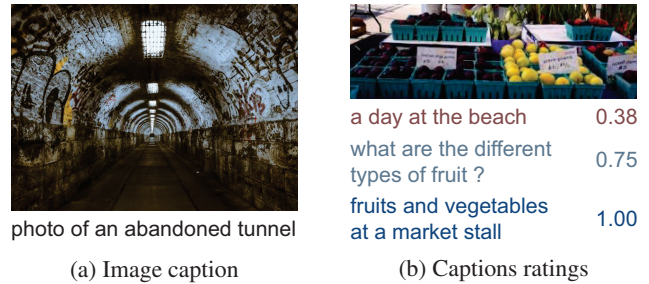


Figure 1: Examples from an image caption dataset (Sharma et al. 2018) and a caption-ratings dataset (Levinboim et al. 2019). (a) Images in the caption dataset are annotated with ground-truth captions written by humans. (b) Captions in the caption ratings dataset are generated by trained models and scored in $[0, 1]$ (0: worst, 1: best) by human raters.

guide the training procedure towards models capable of producing outputs that are highly-rated by human evaluators.

As a result of the need to understand the performance of the current models, human evaluation studies for measuring caption quality are frequently reported in the literature (Sharma et al. 2018; Forbes et al. 2019; Dognin et al. 2019; Zhao et al. 2019). In addition to an aggregate model performance, such human evaluation studies also produce a valuable by-product: a dataset of model-generated image captions with human annotated quality labels, as shown in Figure 1b. We argue that such a by-product, henceforth called a caption ratings dataset, can be successfully used to improve the quality of image captioning models, for several reasons. First, optimizing based on instance-level human judgments of caption quality represent a closer-to-truth objective for image captioning: generating more captions judged as good but fewer ones rated as poor by human raters. Second, while having highly-rated captions as positive examples (*i.e.*, how good captions may look like), a caption ratings dataset also contains captions that are highly-scored by a model but annotated as negative examples (*i.e.*, how model-favored yet bad captions look like), which intuitively should be a useful signal for correcting common model biases. To the best of our knowledge, our work is the first to propose using human

*Work done during internship at Google.

caption ratings directly for training captioning models.

Our goal is to leverage the signals from a pre-collected caption ratings dataset (Levinboim et al. 2019) for training an image captioning model. We propose a method based on policy gradient, where the human ratings are considered as rewards for generating captions (seen as taking actions) in an RL framework. Since the dataset provides ratings only for a small set of images and captions, we do not have a generic reward function for random image-caption pairs. Therefore, it is not straightforward to apply policy gradient method that requires a reward for randomly sampled captions. To address this challenge, we use an off-policy technique and force the network to sample captions for which ratings are available in the dataset. We evaluate the effectiveness of our method using human evaluation studies on the T2 test set used for the Conceptual Captions Challenge¹, using both a similar human evaluation methodology and an additional, multi-dimensional side-by-side human evaluation strategy. Additionally, the human raters in our evaluation study are different from the ones that provided the caption ratings in (Levinboim et al. 2019), thereby ensuring that the results are independent of using a specific human-evaluator pool. The results of our human evaluations indicate that the proposed method improves the image captioning quality, by effectively leveraging *both* the positive and negative signals from the captions ratings dataset.

The main contributions of this paper are the following:

- We propose to train captioning models using human ratings produced during evaluations of previous models.
- We propose an off-policy policy gradient method to cope with the sparsity in available caption ratings.
- We present a set of experiments using human evaluations that demonstrates the effectiveness of our approach.

2 Related Work

There have been multiple attempts to define metrics that evaluate the quality of generated captions. Several studies proposed automatic metrics using ground-truth captions. A few of them are adopted from machine translation community and are based on n -gram matches between ground-truth and generated captions; BLEU (Papineni et al. 2002) and ROUGE (Lin 2004) measures precision and recall based on n -gram matches, respectively, while METEOR (Banerjee and Lavie 2005) incorporates alignments between n -gram matches. In the context of evaluating image caption quality specifically, CIDEr (Vedantam, Zitnick, and Parikh 2015) and SPICE (Anderson et al. 2016) utilize more corpus-level and semantic signals to measure matches between generated and ground-truth captions. Aside from these handcrafted metrics, a recent study proposes to learn an automatic metric from a captioning dataset (Cui et al. 2018), while another uses semantic similarity between object labels identified in the image and the words in the caption (Madhyastha, Wang, and Specia 2019).

To overcome the limitations imposed by the automatic metrics, several studies evaluate their models using hu-

man judgments (Sharma et al. 2018; Zhao et al. 2019; Dognin et al. 2019; Forbes et al. 2019). However, none of them utilizes the human-rated captions in the model evaluations. In this work, we show how one can utilize such human-rated captions for training better captioning models.

MLE with ground-truth captions has been widely adopted as the standard surrogate objective for training (Vinyals et al. 2015; Mun, Cho, and Han 2017; Lu et al. 2017; 2018; Changpinyo et al. 2019). Aside from this main thrust, an additional line of research is concerned with optimizing models that maximize some automatic evaluation metric(s) using RL, in an attempt to bridge the mismatch between the training objective and the evaluation criteria (Rennie et al. 2017; Ding and Soricut 2017; Anderson et al. 2018; Qin et al. 2019). To our knowledge, this is the first study that proposes to optimize test-time scores of human judgment using a dataset generated by a human evaluation process.

Another line of related research is focused on learning from human feedback, which has been actively explored in the field of RL. Some approaches use binary human feedback to update an agent (Knox et al. 2012; Amershi et al. 2014; MacGlashan et al. 2017) whereas approaches with preference-based RL take human feedback as preferences between two action/state trajectories (Akrou, Schoenauer, and Sebag 2012; Wirth, Fürnkranz, and Neumann 2016; Wirth et al. 2017). A common technique adopted in these methods is to learn an estimation model from human feedback to approximate the absent reward function (Knox et al. 2012; Christiano et al. 2017; Ibarz et al. 2018). However, these approaches assume that the models receive human feedback iteratively in a training loop; in contrast, our approach uses the caption ratings in an off-line manner, simply as a pre-existing annotation dataset. As a result, our method focuses on existing examples within the dataset, using an off-policy technique.

3 Methods

3.1 Caption Ratings Dataset

A sample in a caption ratings dataset is comprised of an image I , a machine-generated caption c , and a human judgment for the caption quality $r(c|I) \in \mathbb{R}$. For each image, multiple captions from several candidate models are available, some of which might be rated higher than others. In the setup used in this paper, the low-rated captions serve as negative examples, because human annotators judged them as bad captions (see examples in Figure 1b). $r(c|I)$ is possibly an aggregate of multiple ratings from different raters. Section 4.1 provides more details of the caption ratings dataset that we employ.

We make a few observations that apply not only to image captioning, but more generally to the principle of generating annotations. Although a human-ratings dataset is usually just a by-product of human evaluations for past models, such a dataset can be valuable for improving models (as we show in this paper). There are several advantageous properties of a ratings dataset over traditional supervised-learning datasets. First, obtaining ratings for automatically generated outputs is significantly cheaper than collecting ground-truth labels, because it requires less rater training and less time spent an-

¹<http://www.conceptualcaptions.com/challenge>




Figure 2: Illustration of training settings for three different methods. The 2D boxes represent the space of possible captions, and each dot is a caption with its color corresponding to human ratings. (Blue is high and red is low.) The solid line in each plot indicates the virtual boundary separating low-quality and high-quality captions in terms of human ratings. The color gradation in (b) represents learned rating estimates for captions, while the dashed line is the model’s approximated virtual boundary between low- and high-quality estimates. In (c), $\eta = \frac{p_{\theta}(c'_s|I)}{q(c'_s|I)}$ denotes the importance weight for a sample c'_s .

notating. Moreover, if human evaluation is performed anyway during a model’s development cycle, there is no additional cost associated to using these annotations for further improving the model. In addition to that, it is easy to capture consensus between multiple raters to reduce noise, *e.g.*, by averaging their scores; it is completely non-trivial to achieve a similar effect from multiple ground-truth labels. Last but not least, the examples with a negative rating score provide valuable training signals, as they explicitly penalize the mistakes that appear in model outputs with high-probability; this type of signal is completely lacking in traditional supervised-learning datasets.

3.2 Reinforcing Caption Generator using Ratings

Given a caption ratings dataset \mathcal{D} with triplets $(I, c, r(c|I))$, our objective is to maximize the expected ratings of the output captions $\mathcal{J}(\theta)$, which is given by

$$\begin{aligned} \mathcal{J}(\theta) &= \mathbb{E}_{I \sim p_{\mathcal{D}}(I), c \sim p_{\theta}(c|I)}[r(c|I)] \\ &= \sum_I p_{\mathcal{D}}(I) \sum_c p_{\theta}(c|I) r(c|I), \end{aligned} \quad (1)$$

where $p_{\mathcal{D}}(I)$ is the dataset distribution for I and $p_{\theta}(c|I)$ is the conditional caption distribution estimated by a model parameterized by θ .

Our objective in Eq. (1) exactly aligns with the reward maximization of RL, and therefore we apply the techniques of RL by configuring the captioning model as the agent, the rating scores as the reward, the input images as the states, and the captions as the actions. Specifically, we use a policy gradient method where an approximated policy gradient is computed using Monte-Carlo sampling,

$$\begin{aligned} \nabla_{\theta} \mathcal{J}_{\text{PG}}(\theta) &= \mathbb{E}_{\pi}[(r(c|I) - b) \nabla_{\theta} \ln p_{\theta}(c|I)] \\ &\approx \frac{1}{S} \sum_{s=1}^S (r(c_s|I_s) - b) \nabla_{\theta} \ln p_{\theta}(c_s|I_s), \end{aligned} \quad (2)$$

where \mathbb{E}_{π} represents $\mathbb{E}_{I \sim p_{\mathcal{D}}(I), c \sim p_{\theta}(c|I)}$, I_s and c_s are image and caption sampled from $p_{\mathcal{D}}(I)$ and $p_{\theta}(c|I)$, respectively, and S is the number of samples. In the above equations, we

subtract a baseline b from the rating score $r(c_s|I_s)$ to reduce the variance of the estimator while keeping its original bias.

Although this formulation is straightforward, there remains a critical challenge to apply this technique to our task, since the dataset \mathcal{D} contains only sparse information about $r(c|I)$ and true ratings for most captions are unknown. Eq. (2) requires the rating $r(c_s|I_s)$ for a randomly sampled caption which may not be present in the dataset \mathcal{D} . In the rest of this section, we present two alternative techniques for this challenge, and discuss the advantages of one alternative versus the other.

On-policy policy gradient with rating estimates One approach to address the sparsity of the rating function is to construct a caption quality estimator, while keeping the sampling process on-policy; this is the method adopted in, *e.g.*, (Knox et al. 2012; Christiano et al. 2017; Ibarz et al. 2018). Incidentally, it is also the expressed goal for the effort behind the caption ratings dataset in (Levinboim et al. 2019) that we use in this work.

For this purpose, we train a rating estimator $\tilde{r}(c|I; \phi)$ parameterized by ϕ , by minimizing mean squared error of the true rating scores for the image-caption pairs on the caption ratings dataset. The trained estimator then replaces the true rating function $r(c_s|I_s)$ in Eq. (2) and the estimated policy gradient is now:

$$\nabla_{\theta} \mathcal{J}_{\text{PG}}(\theta) \approx \frac{1}{S} \sum_{s=1}^S (\tilde{r}(c_s|I_s; \phi) - b) \nabla_{\theta} \ln p_{\theta}(c_s|I_s). \quad (3)$$

This technique allows to obtain rating estimates for any image-caption pairs, including ones that are not present in the dataset \mathcal{D} . The training objective with Eq. (3) is now maximizing the expected rating estimate of captions. This approach is effective only if the trained rating estimator generalizes well to unseen images and captions, and it is expected to be effective only to the extent to which the rating estimator performs well over the sampled search space. In our work, we have observed artifacts of the ratings estimator that negatively impact the performance of this method, *e.g.*, severely ill-formed captions for which the caption estimator

had no training signal but assigned high ratings. We report results for this method in Section 4.

Off-policy policy gradient with true ratings This second method takes an orthogonal approach to address the sparsity of the rating function. We modify the sampling process in such a manner that it allows us to directly utilize the true ratings of the dataset (no estimation involved), while ensuring that the training procedure is not influenced by the captions whose true ratings are not available. More precisely, we adopt an off-policy policy gradient technique that uses an alternative distribution $q(c|I)$, instead of the true policy distribution $p_\theta(c|I)$ for sampling. The policy gradient in Eq. (2) is approximated as follows:

$$\begin{aligned}\nabla_\theta \mathcal{J}_{\text{PG}}(\theta) &= \mathbb{E}_\pi[(r(c|I) - b) \nabla_\theta \ln p_\theta(c|I)] \\ &= \mathbb{E}_\beta \left[\frac{p_\theta(c|I)}{q(c|I)} (r(c|I) - b) \nabla_\theta \ln p_\theta(c|I) \right] \\ &\approx \frac{1}{S} \sum_{s=1}^S \frac{p_\theta(c_s|I_s)}{q(c_s|I_s)} (r(c_s|I_s) - b) \nabla_\theta \ln p_\theta(c_s|I_s),\end{aligned}\quad (4)$$

where \mathbb{E}_β represents $\mathbb{E}_{I \sim p_{\mathcal{D}}(I), c \sim q(c|I)}$ with an alternative caption distribution $q(c|I)$, and $\frac{p_\theta(c|I)}{q(c|I)}$ represents the importance weight for sample caption c_s and image I_s . The alternative caption sampling distribution is defined as:

$$q(c|I) = (1 - \epsilon)p_{\mathcal{D}}(c|I) + \epsilon U(c), \quad (5)$$

where $p_{\mathcal{D}}(c|I)$ is the conditional caption distribution in the dataset \mathcal{D} , $U(\cdot)$ is the uniform distribution, and $\epsilon \ll 1$ is a small positive weight assigned to the uniform distribution. In all experiments, we sample a single caption per image in the batch. While captions that are not present in the dataset may still be sampled from $U(c)$, we assign a reward b to these captions, in order to prevent incorrect contributions to the gradient computation. In the policy gradient formulation, examples with reward value b are considered to have no information, and their weight $r(c|I) - b = 0$ cancels out the entire term corresponding to these examples. Note that the off-policy methods enable experience replay, which is repeating previous experiences with known rewards. In this view, this method is viewed as training a captioning model by replaying the experiences in the ratings dataset.

Curriculum learning As our training conditions, we assume the access to both a captioning dataset and a caption ratings dataset. Under a curriculum learning procedure, we first train a model by MLE on the captioning dataset, and then fine-tune the model with the above methods using the caption ratings dataset. To avoid overfitting during fine-tuning, we add the MLE loss on the captioning dataset as a regularization term. Given the caption labeled dataset \mathcal{D}_{IC} and the caption ratings dataset \mathcal{D}_{CR} , the final gradients w.r.t. the parameters are therefore computed as follows:

$$\nabla_\theta \mathcal{J}(\theta) = \alpha \nabla_\theta \mathcal{J}_{\text{PG}}(\theta; \mathcal{D}_{\text{CR}}) + \nabla_\theta \mathcal{J}_{\text{MLE}}(\theta; \mathcal{D}_{\text{IC}}), \quad (6)$$

where \mathcal{J}_{MLE} is the average log-likelihood of ground-truth captions in \mathcal{D}_{IC} , and α is a hyper-parameter that balances the regularization effect.

3.3 Comparing two policy gradient methods

Intuitively, the two policy gradient methods described in this section have strong relationships to MLE, since training signals are based on the gradients of caption log-likelihoods. We illustrate the training settings of MLE and the two proposed methods in Figure 2. In MLE, we train the model using positive captions only and treat all positive captions equally, as illustrated in Figure 2a: the parameters are updated by the gradients of log-likelihoods of ground-truth captions c_{GT} . The on-policy policy gradient method (Eq. (3)) instead computes the gradients of reward-weighted log-likelihoods of sample captions c_s over all possible captions. By sampling from the policy distribution (on-policy), we may sample captions whose true rating scores are not known (not in the dataset). The on-policy method thus approximates the rating function by a rating estimator $\hat{r}(c|I)$, depicted by the background gradient in Figure 2b. However, the mismatch between the true rating function and the estimator (depicted by the gap between solid and dashed lines) can degenerate the quality of the resulting captioning model. On the other hand, the off-policy method focuses on the captions with true rating scores in the dataset, by changing the sampling distribution. In contrast to MLE, where each sample is viewed as equally correct and important, the off-policy method weights each caption by its rating, and therefore includes captions with negative feedback, as illustrated in Figure 2c. Note that, in the off-policy method, the baseline determines the threshold for positive/negative feedback; captions with ratings below the baseline are explicitly penalized, while the others are positively rewarded.

4 Experiments

4.1 Datasets

Image captioning dataset In the experiments, we use Conceptual Captions (Sharma et al. 2018), a large-scale captioning dataset that consists of images crawled from the Internet, with captions derived from corresponding Alt-text labels on the webpages. The training and validation splits have approximately 3.3M and 16K samples, respectively.

Caption ratings dataset In our experiments, we use the Caption-Quality dataset (Levinboim et al. 2019), recently introduced for the purpose of training quality-estimation models for image captions. We re-purpose this data as our caption ratings dataset \mathcal{D}_{CR} . The dataset is divided into training, validation and test splits containing approximately 130K, 7K and 7K rated captions, respectively. Each image has an average of 4.5 captions (generated by different models that underwent evaluation). The captions are individually rated by asking raters the question “*Is this a good caption for the image?*”, with the answers “*NO*” or “*YES*” mapped to a 0 or 1 score, respectively. Each image/caption pair is evaluated by 10 different human raters, and an average rating score per-caption is obtained by quantizing the resulting averages into a total of nine bins $\{0, \frac{1}{8}, \dots, \frac{7}{8}, 1\}$.

Conceptual Captions Challenge T2 dataset To evaluate our models, we run human evaluation studies on the T2 test

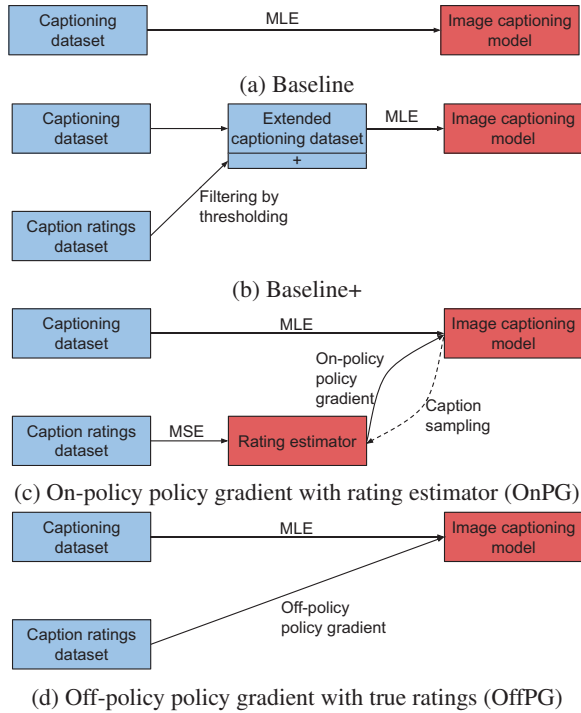


Figure 3: The training procedures for the different methods described. Blue and red boxes represent datasets and models, respectively. MLE is maximum likelihood estimation and MSE means mean-squared error minimization.

Table 1: Questions asked to raters in the two human evaluations. Type ‘S’ evaluation means single-caption rating. Type ‘SxS’ evaluation means side-by-side caption rating.

Metric	Type	Question
Goodness	S	Is this a good caption for the image?
Informativeness	SxS	Which caption provides more useful info for a person who cannot see this image?
Correctness	SxS	Which caption has fewer mistakes?
Fluency	SxS	Which caption has better language quality?

dataset used in the CVPR 2019 Conceptual Captions Challenge². The dataset contains 1K images sampled from the Open Images Dataset (Kuznetsova et al. 2018). Note that the images in the Caption-Quality dataset are also sampled from the Open Images Dataset, but using a disjoint split. So there is no overlap between the caption ratings dataset \mathcal{D}_{CR} we use for training, and the T2 test set we use for evaluations.

4.2 Experimental Settings

Model architecture As the backbone model for image captioning we adopt the architecture described in (Changpinyo et al. 2019), since it provides the highest single-model score in the Conceptual Captions Challenge³. Given an image, we extract two types of visual features: 1) ultra fine-grained semantic features using pretrained network (Juan et

al. 2019) from the entire image and 16 bounding boxes proposed by faster-RCNN (Ren et al. 2015), and 2) label embeddings of objects predicted by Google Cloud Vision API⁴. We use these features with an encoder-decoder Transformer Network (Vaswani et al. 2017) to generate the captions.

In addition, we train a caption rating estimator for the *OnPG* method using the Caption-Quality dataset. The rating estimator extracts the same types of visual features as the captioning model above, and embeds the input caption with a pretrained BERT encoder (Devlin et al. 2018). We concatenate all these features after projecting into a common embedding space and predict the human ratings of the input image/caption pair. To feed the generated captions from the captioning model directly into the rating estimator, we share the vocabulary (but not the token embeddings) between the two models. We fix the pretrained image feature extraction modules in both models during training, as well as the BERT encoder of the rating estimator. The rating estimator achieves a test performance that is close to the one reported (0.519 Spearman correlation) in (Levinboim et al. 2019); however, as we will discuss further, its performance on the Caption-Quality test set does not transfer well to the needs of the *OnPG* method, which needs correct rating estimates for ill-formed captions as well.

Baselines and proposed models We first train an MLE model as our baseline, trained on the Conceptual Captions training split alone. We referred to this model as *Baseline*. For a baseline approach that utilizes (some of) the Caption-Quality data, we merge positively-rated captions from the Caption-Quality training split with the Conceptual Captions examples and finetune the baseline model. We call this model *Baseline+ (t)* , where $t \in [0, 1]$ is the rating threshold for the included positive captions. We train models for two variants, $t \in \{0.5, 0.7\}$, which results in $\sim 72K$ and $\sim 51K$ additional (pseudo-)ground-truth captions, respectively. Note that the *Baseline+ (t)* approaches attempt to make use of the same additional dataset as our two reinforced models, *OnPG* and *OffPG*, but they need to exclude below-threshold captions due to the constraints in MLE.

In addition to the baselines, we train two reinforced models: one based on the on-policy policy gradient method with a rating estimator (*OnPG*), and the other based on the off-policy policy gradient method with the true ratings (*OffPG*). The differences between the methods are shown in Figure 3.

Training details We train *Baseline* using the Adam optimizer (Kingma and Ba 2014) on the training split of the Conceptual dataset for 3M iterations with the batch size of 4,096 and the learning rate of 3.2×10^{-5} . The learning rate is warmed up for 20 epochs and exponentially decayed by a factor of 0.95 every 25 epochs. *Baseline+ (t)* are obtained by fine-tuning *Baseline* on the merged dataset for 1M iterations, with the learning rate of 3.2×10^{-7} and the same decaying factor. For *OnPG*, because its memory footprint is increased significantly due to the additional parameters for the rating estimator, we reduce the batch size for training this model by a 0.25 factor; the value of b in Eq. (2) is set to the

²<http://www.conceptualcaptions.com/winners-and-data>

³As of Sept. 5, 2019.

⁴<http://cloud.google.com/vision>

Table 2: Human evaluation single-caption results: Goodness scores for models (higher is better). Column Δ shows relative improvements over *Baseline*. Note that all score increases of *Baseline*+(t) and *OnPG* are within the error range.

	Average Goodness		Voting Goodness	
		Δ		Δ
Baseline	66.23 \pm 0.60%	—	66.30 \pm 1.49%	—
Baseline+ (0.5)	66.68 \pm 0.61%	-0.45%	66.50 \pm 1.48%	-0.20%
Baseline+ (0.7)	65.83 \pm 0.62%	-0.40%	65.20 \pm 1.51%	-1.10%
OnPG	65.97 \pm 0.61%	-0.26%	66.40 \pm 1.48%	-0.10%
OffPG	68.42\pm0.61%	3.19%	69.70\pm1.46%	3.40%

Table 3: Human evaluation side-by-side comparisons against the baseline. Positive values denote superior performance compared to the baseline. Note that some score increases for *Baseline*+(t) and *OnPG* are within error range.

	Informativeness	Correctness	Fluency
Baseline	—	—	—
Baseline+ (0.5)	1.78 \pm 0.85%	0.18 \pm 0.49%	0.10 \pm 0.28%
Baseline+ (0.7)	0.70 \pm 0.58%	0.68 \pm 0.33%	0.23 \pm 0.15%
OnPG	-0.33 \pm 0.90%	-0.35 \pm 0.62%	0.08 \pm 0.20%
OffPG	7.45\pm1.06%	5.90\pm0.80%	1.69\pm0.31%

moving average of the rating estimates. During *OffPG* training, for each batch, we sample half of the examples from the Conceptual dataset and the other half from Caption-Quality dataset; b is set to the average of the ratings in the dataset.

4.3 Evaluations

We run two sets of human evaluation studies to evaluate the performance of our models and baselines, using the T2 dataset (1K images). For every evaluation, we generate captions using beam search (beam size of 5).

Single-caption evaluation In the first type of evaluation, 6 distinct raters are asked to judge each image caption as good or bad. They are shown the image and caption with the “Goodness” question prompt shown in Table 1. The bad or good rating is translated to 0 or 1, respectively. We measure “average” goodness score as the average of all the ratings over the test set. We also report a “voting”⁵ score which is the average of the binarized score for each caption based on majority voting. Note that both the “average” and “voting” scores are in the range $[0, 1]$, where higher values denote better model performance.

Side-by-side caption evaluation In the other type of evaluation, we measure the relative improvement of a model against the *Baseline* model; Three professional raters are shown the input image and two captions (anonymized and randomly shuffled with respect to their left/right position) side-by-side. One of the captions is from a candidate model and the other always from *Baseline*. We ask for relative judgments on three dimensions – Informativeness, Correctness and Fluency, using their corresponding questions shown in Table 1. Each of these dimensions allows a 5-way choice, shown below together with their corresponding scores:

⁵The “voting” score is the metric reported on the Conceptual Captions Challenge leaderboard.




Figure 4: Rating distribution for the correctness question. Tendencies are similar for the other side-by-side questions.




Figure 5: Results of *OnPG* and *OffPG* in side-by-side human comparisons while varying weight of policy gradient α . Models are tested on 200 samples from T2 dataset.

The left caption is much better	-1.0
The left caption is slightly better	-0.5
The two captions seem equal	0.0
The right caption is slightly better	+0.5
The right caption is much better	+1.0

Each model is evaluated by the average rating scores from 3 distinct raters. As a result, we obtain 3 values for each model in the range $[-1, 1]$, where a negative score means a performance degradation in the given dimension with respect to *Baseline*. For every human evaluation, we report confidence intervals based on bootstrap resampling (Koehn 2004).

4.4 Results

Single-caption evaluation Table 2 shows the goodness scores from the single-caption evaluation. Both “average” and “voting” metrics clearly indicate that *OffPG* significantly improves over *Baseline*, while the other methods achieve only marginal gains, all of which are within the error range. *Baseline*+(t) models use only 1.5% and 2.2% additional data, at $t = 0.7$ and $t = 0.5$, respectively, with insignificant impact. Moreover, these methods only maximize the likelihood of the additional captions, which are already generated with high likelihood by previous models trained on the same dataset, which results in self-reinforcement. In contrast, the policy gradient methods are allowed to utilize the negative feedback to directly penalize incorrect captions. However, *OnPG* fails to improve the quality, most likely because it relies on a noisy caption ratings estimator that fails to generalize well over the large space of possible captions.

Side-by-side evaluations The results from the side-by-side evaluations are shown in Table 3. The *OffPG* method achieves significant improvements on all three different dimensions. This is an important result, considering that





				
Baseline	a view from the garden	runners in the half marathon	this is one of the most beautiful flowers in the world .	person in the water taxi
Baseline+ (0.5)	looking up at the castle (0.8, 0.2, 0.0)	person running in the half marathon (-0.5, 0.0, 0.0)	this is what the flower looks like . (0.2, 0.2, 0.0)	person in the water taxi (0.0, 0.0, 0.0)
Baseline+ (0.7)	a view from the garden (0.0, 0.0, 0.0)	runners in the half marathon (0.0, 0.0, 0.0)	this is what the flower looks like . (0.3, 0.0, 0.0)	person in the water taxi (0.0, 0.0, 0.0)
OnPG	looking up to the castle (1.0, 0.5, 0.0)	person running in a race (-0.5, -0.5, 0.0)	another picture of the beautiful flowers (0.3, 0.2, 0.0)	person in the water taxi (0.0, 0.0, 0.0)
OffPG	the entrance to the castle (1.0, 0.5, 0.0)	runners in the marathon start (1.0, 1.0, 0.0)	a close up of the flowers . (0.8, 0.5, 0.0)	athlete in action in the water (1.0, 1.0, 0.0)

Figure 6: Qualitative examples of generated captions. Numbers represent informativeness, correctness and fluency scores (rated by comparing against those generated by *Baseline*).

we trained the model using a caption ratings dataset that contains single-scalar scores for generic ‘goodness’ (as opposed to the well-defined dimensions along which the *OffPG* method scores have improved). These results demonstrate that the single-caption ‘goodness’ ratings encapsulate a signal for all these dimensions into its scalar value. Note that we observe the same tendency consistently under a variety of hyperparameter settings in our internal experiments.

Figure 4 highlights the way in which the *OffPG* method achieves its superiority over the *Baseline* model, compared to the other alternative models (using the ‘Correctness’ scores). For instance, over 75% of the captions for both *Baseline*+(t) models receive a 0.0 score (equal quality), and more than half of them are exactly identical to their corresponding *Baseline* captions. In contrast, *OffPG* makes a strong impact by explicitly penalizing the captions with negative feedback: less than 16% captions are identical to the corresponding *Baseline* captions. Moreover, we observe a large portion of captions with scores of 1.0 in favor of *OffPG*, indicating that many captions are significantly enhanced. We observe similar trends in all the three metrics.

On-policy vs. off-policy performance We compare the *OnPG* and *OffPG* methods in more depth, by performing ablation experiments for the α hyper-parameter (the weight for the policy gradient). Figure 5 shows the results of these ablation experiments, for which we performed side-by-side comparisons over a 200-image subset from the T2 dataset. The results indicate that a very small α limits the impact of the additional signal for both models, since the regularization effect from the original loss term becomes too strong. By allowing updates using policy gradient with a larger α value, *OffPG* improves the performances along all three dimensions, whereas the performance of *OnPG* starts degrading at higher α values. At $\alpha = 100$, *OnPG* drastically suffers from mode collapse and ends up generating a single caption for every image. This mode collapse is a result of poor generalization of the rating estimator: the collapsed captions are structurally ill-formed (e.g., an empty string, or a string with simply a period ‘.’), but they receive high rating estimates (> 0.9) from the estimator. Although we can (and did) introduce some heuristics to avoid some of these failure cases in the estimator, we observe that *OnPG* training would con-

tinue to suffer from the estimator failing to generalize well over the vast space of possible captions. This observation is similar to the mode collapsing phenomenon seen when training generative adversarial networks (GANs), but even more severe as the estimator in *OnPG* is fixed (unlike the discriminators in GANs which are trained simultaneously).

Another drawback of *OnPG* is that it increases the computational complexity significantly during training. In terms of the memory usage, the rating estimator introduces 65% additional parameters, and uses more than double the memory for gradient computation compared to the other models. Also, the sequential caption sampling in *OnPG* slows down the training procedure, by breaking the parallelism in the Transformer computations, in addition to the time complexity incurred by the rating estimator. Empirically, *OnPG* is over 10 times slower than the others in processing the same number of examples in training. In contrast, the time and space complexities of *OffPG* remain the same as *Baseline* and *Baseline*+(t), since the only difference is the use of scalar weights ($r(c|I)$ and η) to gradients of each caption likelihood ($\nabla_{\theta} \ln p_{\theta}(c|I)$), as shown in Figure 2.

Qualitative results Figure 6 presents some qualitative example outputs for our models, showcasing the effectiveness of the *OffPG* method. We observe that the *OffPG* model is often successful at correcting arbitrary qualifiers present in the baseline outputs (e.g., ‘half marathon’ and ‘most beautiful’ in the second and third examples, respectively).

5 Conclusion

In this paper, we describe how to train an improved captioning model by using a caption ratings dataset, which is often a natural by-product in the development process of image captioning models. We show that an off-policy RL technique with an alternative sampling distribution successfully deals with the sparsity of information about the rating function, while an on-policy method has difficulties in obtaining an improved model, due to generalization issues of the ratings estimator. While this conclusion may not be definitive, it is definitely an important result, and it also opens up additional lines of inquiry along the relative merits of these RL techniques.

References

- Akrou, R.; Schoenauer, M.; and Sebag, M. 2012. April: Active preference learning-based reinforcement learning. In *ECML-KDD*, 116–131. Springer.
- Amershi, S.; Cakmak, M.; Knox, W. B.; and Kulesza, T. 2014. Power to the people: The role of humans in interactive machine learning. *AI Magazine* 35(4):105–120.
- Anderson, P.; Fernando, B.; Johnson, M.; and Gould, S. 2016. SPICE: Semantic propositional image caption evaluation. In *ECCV*, 382–398. Springer.
- Anderson, P.; He, X.; Buehler, C.; Teney, D.; Johnson, M.; Gould, S.; and Zhang, L. 2018. Bottom-up and top-down attention for image captioning and visual question answering. In *CVPR*, 6077–6086.
- Banerjee, S., and Lavie, A. 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*, 65–72.
- Changpinyo, S.; Pang, B.; Sharma, P.; and Soricut, R. 2019. Decoupled box proposal and featurization with ultrafine-grained semantic labels improve image captioning and visual question answering. In *EMNLP-IJCNLP*.
- Christiano, P. F.; Leike, J.; Brown, T.; Martic, M.; Legg, S.; and Amodei, D. 2017. Deep reinforcement learning from human preferences. In *NIPS*, 4299–4307.
- Cui, Y.; Yang, G.; Veit, A.; Huang, X.; and Belongie, S. 2018. Learning to evaluate image captioning. In *CVPR*.
- Devlin, J.; Chang, M.-W.; Lee, K.; and Toutanova, K. 2018. BERT: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*.
- Ding, N., and Soricut, R. 2017. Cold-start reinforcement learning with softmax policy gradients. In *NIPS*.
- Dognin, P.; Melnyk, I.; Mroueh, Y.; Ross, J.; and Sercu, T. 2019. Adversarial semantic alignment for improved image captions. In *CVPR*, 10463–10471.
- Forbes, M.; Käser-Chen, C.; Sharma, P.; and Belongie, S. 2019. Neural naturalist: Generating fine-grained image comparisons. In *EMNLP-IJCNLP*.
- Ibarz, B.; Leike, J.; Pohlen, T.; Irving, G.; Legg, S.; and Amodei, D. 2018. Reward learning from human preferences and demonstrations in atari. In *NeurIPS*, 8011–8023.
- Juan, D.-C.; Lu, C.-T.; Li, Z.; Peng, F.; Timofeev, A.; Chen, Y.-T.; Gao, Y.; Duerig, T.; Tomkins, A.; and Ravi, S. 2019. Graph-rise: Graph-regularized image semantic embedding. *arXiv preprint arXiv:1902.10814*.
- Kingma, D. P., and Ba, J. 2014. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*.
- Knox, W. B.; Glass, B. D.; Love, B. C.; Maddox, W. T.; and Stone, P. 2012. How humans teach agents. *International Journal of Social Robotics* 4(4):409–421.
- Koehn, P. 2004. Statistical significance tests for machine translation evaluation. In *EMNLP*, 388–395.
- Kuznetsova, A.; Rom, H.; Alldrin, N.; Uijlings, J.; Krasin, I.; Pont-Tuset, J.; Kamali, S.; Popov, S.; Mallocci, M.; Duerig, T.; et al. 2018. The open images dataset v4: Unified image classification, object detection, and visual relationship detection at scale. *arXiv preprint arXiv:1811.00982*.
- Levinboim, T.; Thapliyal, A.; Sharma, P.; and Soricut, R. 2019. Quality estimation for image captions based on large-scale human evaluations. *arXiv preprint arXiv*.
- Lin, C.-Y. 2004. ROUGE: A package for automatic evaluation of summaries. In *Text summarization branches out*.
- Lu, J.; Xiong, C.; Parikh, D.; and Socher, R. 2017. Knowing when to look: Adaptive attention via a visual sentinel for image captioning. In *CVPR*, 375–383.
- Lu, J.; Yang, J.; Batra, D.; and Parikh, D. 2018. Neural baby talk. In *CVPR*, 7219–7228.
- MacGlashan, J.; Ho, M. K.; Loftin, R.; Peng, B.; Wang, G.; Roberts, D. L.; Taylor, M. E.; and Littman, M. L. 2017. Interactive learning from policy-dependent human feedback. In *ICML*, 2285–2294. JMLR. org.
- Madhyastha, P.; Wang, J.; and Specia, L. 2019. VIFIDEL: Evaluating the visual fidelity of image descriptions. In *ACL*.
- Mun, J.; Cho, M.; and Han, B. 2017. Text-guided attention model for image captioning. In *AAAI*.
- Papineni, K.; Roukos, S.; Ward, T.; and Zhu, W.-J. 2002. BLEU: a method for automatic evaluation of machine translation. In *ACL*, 311–318.
- Qin, Y.; Du, J.; Zhang, Y.; and Lu, H. 2019. Look back and predict forward in image captioning. In *CVPR*, 8367–8375.
- Ren, S.; He, K.; Girshick, R.; and Sun, J. 2015. Faster r-cnn: Towards real-time object detection with region proposal networks. In *NIPS*, 91–99.
- Rennie, S. J.; Marcheret, E.; Mroueh, Y.; Ross, J.; and Goel, V. 2017. Self-critical sequence training for image captioning. In *CVPR*, 7008–7024.
- Sharma, P.; Ding, N.; Goodman, S.; and Soricut, R. 2018. Conceptual captions: A cleaned, hypernymed, image alt-text dataset for automatic image captioning. In *ACL*, 2556–2565.
- Vaswani, A.; Shazeer, N.; Parmar, N.; Uszkoreit, J.; Jones, L.; Gomez, A. N.; Kaiser, Ł.; and Polosukhin, I. 2017. Attention is all you need. In *NIPS*.
- Vedantam, R.; Zitnick, C. L.; and Parikh, D. 2015. CIDEr: Consensus-based image description evaluation. In *CVPR*.
- Vinyals, O.; Toshev, A.; Bengio, S.; and Erhan, D. 2015. Show and tell: A neural image caption generator. In *CVPR*.
- Wirth, C.; Akrou, R.; Neumann, G.; and Fürnkranz, J. 2017. A survey of preference-based reinforcement learning methods. *The Journal of Machine Learning Research* 18(1):4945–4990.
- Wirth, C.; Fürnkranz, J.; and Neumann, G. 2016. Model-free preference-based reinforcement learning. In *AAAI*.
- Zhao, S.; Sharma, P.; Levinboim, T.; and Soricut, R. 2019. Informative image captioning with external sources of information. In *ACL*.