

Comparison of Diverse Decoding Methods from Conditional Language Models

Daphne Ippolito* Reno Kriz* Maria Kustikova João Sedoc Chris Callison-Burch

*Authors contributed equally

University of Pennsylvania

{daphnei, rekriz, mkust, joao, ccb}@seas.upenn.edu

Abstract

While conditional language models have greatly improved in their ability to output high-quality natural language, many NLP applications benefit from being able to generate a diverse set of candidate sequences. Diverse decoding strategies aim to, within a given-sized candidate list, cover as much of the space of high-quality outputs as possible, leading to improvements for tasks that re-rank and combine candidate outputs. Standard decoding methods, such as beam search, optimize for generating high likelihood sequences rather than diverse ones, though recent work has focused on increasing diversity in these methods. In this work, we perform an extensive survey of decoding-time strategies for generating diverse outputs from conditional language models. We also show how diversity can be improved without sacrificing quality by over-sampling additional candidates, then filtering to the desired number.

1 Introduction

Conditional neural language models, which train a neural net to map from one sequence to another, have had enormous success in natural language processing tasks such as machine translation (Sutskever et al., 2014; Luong et al., 2015), text summarization (Nallapati et al., 2016), and dialog systems (Vinyals and Le, 2015). These models output a probability distribution over the next token in the output sequence given the input and the previously predicted tokens. Since computing the overall most likely output sequence is intractable, early work in neural machine translation found that beam search is an effective strategy to heuristically sample sufficiently likely sequences from these probabilistic models (Sutskever et al., 2014). However, for more open-ended tasks, beam search is ill-suited to generating a set of diverse candidate sequences; this is because candidates

Beam Search

A bus is stopped at a bus stop.
A bus is parked at a bus stop.
A bus stopped at a bus stop in a city.
A bus stopped at a bus stop at a bus stop.
A bus that is parked in front of a building.

Random Sampling

A bus parked at a bus stop at a bus stop.
There is a bus that is at the station.
A man standing by a bus in a city.
A bus pulling away from the train station.
A bus stopped at a stop on the sunny day.



Figure 1: An image with the top five captions from standard beam search and from random sampling. Note the latter set is more diverse but lower quality.

outputted from a large-scale beam search often only differ by punctuation and minor morphological variations (Li and Jurafsky, 2016).

The term “diversity” has been defined in a variety of ways in the literature, with some using it as a synonym for sentence interestingness or unlikelihood (Hashimoto et al., 2019), and others considering it a measure of how different two or more sentences are from each other (Vijayakumar et al., 2016; Gimpel et al., 2013). We take the latter approach, and define diversity as the ability of a generative method to create a set of possible outputs that are each valid given the input, but vary as widely as possible in terms of word choice, topic, and meaning.

There are a number of reasons why it is desirable to produce a set of diverse candidate outputs for a given input. For example, in collaborative story generation, the system makes suggestions to a user for what they should write next (Clark et al., 2018). In these settings, it would be beneficial to show the user multiple different ways to continue their story. In image captioning, any one sentence-long caption is probably missing some information about the image. Krause et al. (2017) show how a set of diverse sentence-length image captions can be transformed into an entire paragraph about the image. Lastly, in applica-

tions that involve reranking candidate sequences, the reranking algorithms are more effective when the input sequences are diverse. Reranking diverse candidates has been shown to improve results in both open dialog and machine translation (Li et al., 2016a; Li and Jurafsky, 2016; Gimpel et al., 2013). Furthermore, in open-ended dialog, the use of reranking to personalize a model’s responses for each user is a promising research direction (Choudhary et al., 2017).

With these sorts of applications in mind, a variety of alternatives and extensions to beam search have been proposed which seek to produce a set of diverse candidate responses instead of a single high likelihood one (Li et al., 2016a; Vijayakumar et al., 2016; Kulikov et al., 2018; Tam et al., 2019). Many of these approaches show marked improvement in diversity over standard beam search across a variety of generative tasks. However, there has been little attempt to compare and evaluate these strategies against each other on any single task.

In this paper, we survey existing methods for promoting diversity in order to systematically investigate the relationship between diversity and perceived quality of output sequences of conditional language models. In addition to standard beam search and greedy random sampling, we compare several recently proposed modifications to both methods. In addition, we propose the use of over-sampling followed by post-decoding clustering to remove similar sequences.

The main contributions of this paper can be summarized as follows:

- A detailed comparison of existing diverse decoding strategies on two tasks: open-ended dialog and image captioning, and recommendations for a diverse decoding strategy.
- A novel clustering-based algorithm that can be used on the results of any decoding strategy to increase quality and diversity.¹

2 Standard Decoding Methods

Conditional language models, which have wide applications across machine translation, text simplification, conversational agents, and more, generally consist of an encoder, which transforms some input \mathbf{x} into a fixed-size latent representation, and a decoder which transforms these representations in order to output a conditional

¹Code can be found at <https://github.com/rekriz11/DeDiv>.

probability of each word in the target sequence given the previous words and the input. Let $z_t = f(y_1, \dots, y_{t-1}, \mathbf{x})$ represent the output of an encoder-decoder model given input \mathbf{x} and the sequence of tokens predicted so far, y_1, \dots, y_{t-1} , which for notational simplicity we write as $y_{<t}$. The output $z_t \in \mathbb{R}^V$ (where V is the cardinality of the enumerated vocabulary \mathcal{V})

The probability distribution over the next possible token being word $w_i \in \mathcal{V}$ is the softmax:

$$P(y_t = w_i | y_{<t}, \mathbf{x}) = \frac{\exp(z_{t,i})}{\sum_{j=1}^V \exp(z_{t,j})}$$

$$\forall i \in \{1, \dots, V\}$$

Most decoding strategies strive to find the most likely overall sequence, i.e. pick a $\hat{\mathbf{y}}$ such that:

$$\hat{\mathbf{y}} = \arg \max_{\mathbf{y}} P(\mathbf{y} | \mathbf{x}) = \arg \max_{\mathbf{y}} \prod_{t=1}^N P(y_t | y_{<t}, \mathbf{x})$$

Unlike Markovian processes, no sub-exponential algorithm exists to find the optimal decoded sequence, and thus we instead use approximations.

Arg-max The simplest approach to decoding a likely sequence is to greedily select a word at each timestep:

$$\hat{y}_t = \arg \max_{y_t} P(y_t | y_{<t}, \mathbf{x})$$

However, because this deterministic approach typically yields repetitive and short output sequences, and does not permit generating multiple samples, it is rarely used in language modelling.

Random Sampling Another option is to randomly sample from the model’s distribution at every timestep. Often, a temperature parameter T is added to control the entropy of the distribution before sampling.

$$P(y_t = w_i | y_{<t}, \mathbf{x}) = \frac{\exp(z_{t,i}/T)}{\sum_{j=1}^V \exp(z_{t,j}/T)}$$

$$\forall i \in \{1, \dots, V\}$$

$$\hat{y}_t \sim y_t$$

Choosing a temperature greater than one causes outputs to look increasingly more random, while bringing the temperature less than zero causes sequences to increasingly resemble greedy sampling.

Recently, top- s random sampling has been proposed as an alternative to using temperature. Sampling is restricted to the s most likely tokens

at each step (Fan et al., 2018; Radford et al., 2019). We find that top- s random sampling’s hard-restriction on generating low probability words is more effective at controlling the stochasticity of sampled sequences than sampling with temperature.

Beam Search Beam search approximates finding the most likely sequence by performing breadth-first search over a restricted search space. At every step of decoding, the method keeps track of b partial hypotheses. The next set of partial hypotheses are chosen by expanding every path from the existing set of b hypotheses, and then choosing the b with the highest scores. Most commonly, the log-likelihood of the partial sequence is used as the scoring function. We use this as our baseline.²

Since beam search only explores a limited portion of the overall search space, it tends to yield multiple variants of the same high-likelihood sequence, sequences that often only differ in punctuation and minor morphological changes (Li and Jurafsky, 2016). Therefore, standard beam search is not ideal for producing diverse outputs.

3 Extensions to Beam Search

In this section, we will discuss a variety of methods that have been developed recently to eliminate redundancy during decoding and generate a wider range of candidate outputs.

Noisy Parallel Approximate Decoding Introduced by Cho (2016), NPAD is a technique that can be applied to any decoding setting. The main idea is that diversity can be achieved more naturally by taking advantage of the continuous manifold on which neural nets embed language. Instead of encouraging diversity by manipulating the probabilities outputted from the model, diverse outputs are instead produced by adding small amounts of noise to the hidden state of the decoder at each step. The noise is randomly sampled from a normal distribution. The variance is gradually annealed from a starting σ_0 to 0 as decoding progresses (that is $\sigma_t = \frac{\sigma_0}{t}$) under the reasoning that uncertainty is greatest at the beginning of decoding. NPAD can be used in conjunction with any decoding strategy; following the best results from the original paper, we show results using NPAD with beam search.

Extensions to NPAD have sought to learn the direction in which to manipulate the hidden states

²We present the beam search algorithm in the appendix.

using an arbitrary decoding objective (Gu et al., 2017). Since such objectives can be highly domain-specific, we do not evaluate this method.

Top- g Capping In beam search, it is often the case that one hypothesis h is assigned a much higher probability than all other hypotheses, causing all hypotheses in the next step to have h as their parent. Following Li and Jurafsky (2016) and Li et al. (2016b), we add an additional constraint to standard beam search to encourage the model to choose options from diverse candidates. At each step t , current hypotheses are grouped according to the parental hypothesis they come from. After grouping candidates, only the top g from each grouping are considered. The resulting $b \times g$ candidates are ranked, and the top b are selected as hypotheses for the next beam step.

Hamming Diversity Reward Vijayakumar et al. (2016) proposes adding an additional diversity-promoting term, θ , to the log-likelihood before reranking. This term measures how different a candidate hypothesis $c_{\leq t}^{(i)}$ is from the partial hypotheses selected in the previous step. Let $\mathcal{H}_{t-1} = \{c_{\leq t-1}^{(1)}, \dots, c_{\leq t-1}^{(b)}\}$ be these partial hypotheses. Then the beam search scoring function for the i th candidate at timestep t becomes:

$$\text{score}(c_{\leq t}^{(i)}) = \sum_{j=1}^t (\log P(c_j^{(i)} | c_{< j}^{(i)}, \mathbf{x})) + \lambda \theta(c_{\leq t}^{(i)}, \mathcal{H}_{t-1})$$

where λ is a tunable hyperparameter. Vijayakumar et al. (2016) try a variety of definitions for θ , including embedding diversity and n -gram diversity, but they find that Hamming distance, the number of tokens in the candidate sequence which exist in the previously selected partial hypotheses, is most effective. We take the negative of the Hamming distance as θ .

Iterative Beam Search In an attempt to improve the size of the search space explored without sacrificing runtime, Kulikov et al. (2018) propose an iterative beam search method. Beam search is run many times, where the states explored by subsequent beam searches are restricted based on the intermediate states explored by previous iterations. Formally, we can define the set of all partial hypotheses for beam search instance i at time step t as $\mathcal{H}_t^{(i)}$. From here, the search space explored by beam search instance i can be expressed as $S_i = \cup_{t=1}^T \mathcal{H}_t^{(i)}$. The i th beam search is pre-

Method	Description	Method	Description
Random Sampling	Standard decoding mechanism, greedily samples a token from the distribution at each time step.	Random Sampling with Temperature	Before sampling, modify entropy of predicted distribution.
Top- s Random Sampling (Fan et al., 2018)	Restrict sampling to the s -most likely words in the distribution. (story generation)	Beam Search	Standard decoding mechanism, keeps the top b partial hypotheses at every time step. (machine translation)
NPAD Beam Search (Cho, 2016)	Add random noise to the hidden state of the decoder at each time step. (machine translation)	Top- g Capping Beam Search (Li and Jurafsky, 2016)	Only consider the top c hypotheses from each parent hypothesis at each time step. (machine translation, dialog)
Hamming Diversity Beam Search (Vijayakumar et al., 2016)	Penalize new hypotheses that have many of the same tokens as existing partial hypotheses. (image captioning)	Iterative Beam Search (Kulikov et al., 2018)	Run beam search several times, preventing later iterations from generating intermediate states already explored. (dialog)
Clustered Beam Search (Tam et al., 2019)	Initially consider more hypotheses at each time step, and then cluster similar hypotheses together. (dialog)	Post-Decoding Clustering (Ours)	Sample a large number of candidates, and then cluster similar outputs together.

Table 1: Brief high-level descriptions of each decoding method we consider in this paper. In parentheses we give the applications on which the technique was originally applied.

vented from generating any partial hypothesis that has previously been generated, that is, any hypothesis found in $S_{<i} = \cup_{i'=0}^{i-1} S_{i'}$.

The authors also attempt a soft inclusion criterion, where any states within ϵ Hamming distance from a previously explored state are also excluded. During the experimentation of Kulikov et al. (2018), however, the soft-inclusion was found to not be beneficial; thus, we only restrict exact matches of previous states in our implementation. In practice, this means after the first beam search instance runs as normal, the first step of the second beam search instance will contain the $b+1$ to $2b$ -most likely starting tokens; this pattern holds for the third beam search instance, and so on.

Clustered Beam Search Most recently, Tam et al. (2019) proposed a clustering-based beam search method to help condense and remove meaningless responses from chatbots. Specifically, at each decoding step t , this method initially considers the top $2*b$ candidates. From there, each candidate sequence is embedded³, and the embeddings are clustered into c clusters using K -means. Finally, we take the top $\frac{b}{c}$ candidates from each cluster. Note that in the case any clusters have size less than $\frac{b}{c}$, we then include the highest-ranked candidates not found after clustering.

³We follow Tam et al. (2019) and used averaged GloVe word embeddings (Pennington et al., 2014).

4 Clustering Post-Decoding (PDC)

In the previous section, we discuss several diversity-promoting methods that can be applied during the decoding process. However, it is also possible to encourage additional diversity post-hoc. On the task of sentence simplification, after decoding using a large-scale diversity-promoting beam search (beam size 100), Kriz et al. (2019) then clustered similar sentences together to further increase the variety of simplifications from which to choose. Document embeddings generated via Paragraph Vector (Le and Mikolov, 2014) were used as the sentence embeddings with which to perform K -means.

In this work, we extend this post-decoding clustering idea in three key ways. First, we make use of sentence-level embeddings which leverage the pre-trained language representations from the Bidirectional Encoder Representations from Transformers (BERT) (Devlin et al., 2018).⁴ Second, after clustering, Kriz et al. (2019) took the sentence closest to the centroid of each cluster as the representative candidate; we instead choose the highest-ranked candidate (according to log-likelihood) from each cluster to ensure the best candidates are still selected. Finally, after performing standard K -means clustering, we found that it was often the case that some clusters contained large numbers of good candidates, while others contained very few candidates that are also

⁴BERT sentence-level embeddings were obtained using <https://github.com/hanxiao/bert-as-service>.

either ungrammatical or otherwise inferior. Thus, in our implementation, we remove clusters containing two or fewer sentences, and then sample a second candidate from each of the remaining clusters, prioritizing selecting candidates from larger clusters first.

5 Experimental Setup

We evaluate the decoding strategies described in the previous sections under the following settings. For each of the published beam search algorithms, we choose the hyperparameters that were found to be best in the original publications.

RS	Random sampling with temp = 0.5, 0.7, 1.0, or 1.0 with top-10 capping.
Standard BS	Standard beam search
Top5Cap BS	Top- g capping with $g = 3$
Iter5 BS	Iterative beam search with 5 iterations
HamDiv0.8 BS	Hamming Diversity with $\lambda = 0.8$
Cluster5 BS	Clustered beam search with 5 clusters
NPAD0.3 BS	Noisy Decoding with $\sigma_0 = 0.3$

For random sampling, we sample 10 outputs, and with beam-search based methods, we use a beam size of 10 to generate 10 outputs. In addition, we show results from oversampling then filtering. We use a beam size of 100 or generate 100 samples through random sampling, and then we select 10 from the 100, either through post-decoding clustering (PDC) or by taking the 10 candidates with highest likelihood.

We examine these decoding strategies on two tasks: open ended dialog and image captioning. For each task, we evaluate both the quality and diversity of the 10 outputs from each strategy.

5.1 Open-ended Dialog Task

In the dialog domain, we use an LSTM-based sequence-to-sequence (Seq2Seq) model implemented in the OpenNMT framework (Klein et al., 2017). We match the model architecture and training data of Baheti et al. (2018). The Seq2Seq model has four layers each in the encoder and decoder, with hidden size 1000, and was trained on a cleaned version of OpenSubtitles (Tiedemann, 2009) to predict the next utterance given the previous one.

Evaluation is performed on 100 prompts from the Cornell Movie Dialog Corpus (Danescu-Niculescu-Mizil and Lee, 2011). These prompts are a subset of the 1000 prompts used in Baheti et al. (2018), which were filtered using item response theory for discriminative power.

We report perplexity (PpL), averaged over *all* the top 10 outputs for each example.⁵ Since the quality of open-ended dialog is notoriously difficult to evaluate automatically, we ran a human evaluation task on Amazon Mechanical Turk where annotators were shown a prompt and 5 potential responses generated by any of our decoding methods. Evaluators were asked to provide binary ratings on fluency, adequacy, and interestingness for each response. Overall, we collected 3 human judgments for each of the top ten responses for each of our decoding methods; in other words, we collected 3,000 judgments per method.⁶

5.2 Image Captioning Task

For image captioning, we use a state-of-the-art model introduced in Anderson et al. (2018). We take advantage of Luo (2017)’s open-source implementation and released model parameters trained on MSCOCO (Lin et al., 2014). We evaluate on a test set containing 5000 images.

We report Semantic Propositional Image Caption Evaluation (SPICE) scores, an automatic evaluation metric that has been shown to correlate well with human judgments of quality (Anderson et al., 2016). SPICE measures how well the semantic scene graph induced by the proposed caption matches one induced by the ground truth. In addition to computing SPICE on the top-scoring caption (SPICE@1), we follow Vijayakumar et al. (2016) in reporting Oracle SPICE@10 scores. This is done to show the upper bound on the potential impact diversity can have. We also compute the mean SPICE score across all of the candidate captions for an image. Unlike SPICE@1 and SPICE@10, this metric shows the overall quality of *all* of the candidate captions, which is useful to know for applications that combine diverse candidate output sequences (Krause et al., 2017).

5.3 Evaluating Diversity

To measure the diversity across the generated candidate sequences for a given input, we report **Dist-k**, the total number of distinct k-grams divided by the total number of produced tokens in all of the candidate responses for a prompt (Li et al., 2016a). We report Dist-2 and Dist-4 averaged over the prompts in the test set.

⁵This differs from existing work which computes perplexity over only the top output for each example. For our task we are interested in the quality of all of the generated responses.

⁶The full instructions shown on AMT are in the appendix.

Method	Fluency	Adequacy	Interestingness	Ppl	Dist-1	Dist-2	Ent-2	Ent-4
Reference	0.795	0.732	0.636	–	–	–	–	–
RS 0.7 (sample 10)	0.758	0.399	0.388	35.98	0.63	0.80	4.08	3.84
RS 1.0 (sample10)	0.550	0.303	0.386†	67.99	0.74	0.87	4.35	4.08
RS 1.0,top10 (sample 10)	0.745†	0.418	0.387†	10.33	0.60	0.80	4.12	3.91
Standard BS (10 beams)	0.950	0.621	0.336	4.01	0.37	0.45	3.16	3.01
Top3Cap BS (10 beams)	0.942†	0.603	0.346	4.03	0.37	0.46	3.17	3.03
Iter5 BS (10 beams)	0.903	0.520	0.335	5.42	0.62	0.74	3.68	3.25
HamDiv0.8 BS (10 beams)	0.923	0.599	0.366†	4.56	0.33	0.37	3.08	3.00
Cluster5 BS (10 beams)	0.936	0.582	0.381	4.23	0.39	0.46	3.24	3.06
NPAD0.3 BS (10 beams)	0.942†	0.604†	0.335	4.05	0.36	0.44	3.13	2.99
RS 1.0,top10 (sample 100, rank)	0.922	0.548	0.347	5.10	0.52	0.68	3.54	3.18
RS 1.0,top10 (sample 100, PDC)	0.852	0.494	0.372	6.96	0.63	0.76	3.74	3.27
Standard BS (100 beams, rank)	0.964	0.611	0.332†	4.01	0.44	0.61	3.33	3.05
Standard BS (100 beams, PDC)	0.944	0.599	0.346	4.42	0.57	0.70	3.59	3.21

Table 2: Results on 100 dialog prompts. The first row shows the mean human ratings of the single reference response available for each prompt. The next three rows show results for random sampling, with 10 samples drawn per prompt. The next six rows are variants of beam search using beam size 10. The last four rows use random sampling or standard beam search to generate 100 outputs, then filter down to 10 outputs either through ranking by log-likelihood or by performing post-decoding clustering (PDC). In each section, the highest value is bolded, and statistical ties are marked †.

Method	SPICE			Dist-1	Dist-2	Ent-2	Ent-4
	Mean	@1	@10				
RS 0.7 (sample10)	0.170	0.192	0.278	0.31	0.52	3.67	4.00
RS 1.0 (sample10)	0.133	0.167	0.247	0.44	0.71	4.17	4.26
RS 1.0,top10 (sample10)	0.159	0.183	0.272	0.33	0.59	3.90	4.17
Standard BS (10 beams)	0.194	0.193	0.283	0.18	0.26	2.94	3.18
Top3Cap BS (10 beams)	0.195	0.196	0.282	0.17	0.26	2.93	3.17
HamDiv0.8 BS (10 beams)	0.194	0.194	0.282	0.18	0.27	2.98	3.19
Cluster5 BS (10 beams)	0.191	0.194	0.285	0.19	0.28	3.04	3.25
NPAD0.3 BS (10 beams)	0.191	0.192	0.280	0.18	0.26	2.94	3.17
RS 1.0,top10 (sample100, rank)	0.182	0.188	0.284	0.25	0.41	3.31	3.64
RS 1.0,top10 (sample100, PDC)	0.169	0.188	0.282	0.31	0.52	3.62	3.91
Standard BS (100 beams, rank)	0.188	0.190	0.279	0.20	0.31	3.04	3.32
Standard BS (100 beams, PDC)	0.186	0.192	0.288	0.24	0.38	3.25	3.57

Table 3: Image captioning results for selected random sampling and beam search methods. SPICE@1 measures the SPICE score of the most likely caption. SPICE@10 is the maximum score across the 10 candidates generated by each method. Mean SPICE is the mean score over all 10 candidates. In each section, the best value is bolded.

A limitation of Dist- k is that all k -grams that appear at least once are weighted the same, ignoring the fact that infrequent k -grams contribute more to diversity than frequent ones. Zhang et al. (2018) instead propose an entropy metric, **Ent-k**, defined as:

$$Ent-k = \frac{-1}{\sum_{w \in S} F(w)} \sum_{w \in S} F(w) \log \frac{F(w)}{\sum_{w' \in S} F(w')}$$

where S is the set of all k -grams that appear in candidate responses for an example, and $F(w)$ denotes the frequency of w in the candidate responses.

6 Results

We report results on dialog systems and image captioning in Tables 2 and 3, respectively. As expected, random sampling-based approaches yield outputs with greater diversity but worse quality than beam search-based approaches. Over-sampling then filtering increases the quality of outputs while still ensuring high diversity. In the following sections, we discuss the diversity-quality tradeoff, and then delve further into the results for each method group.

6.1 The Quality Diversity Tradeoff

The goal of diverse decoding strategies is to generate high-quality candidate sequences which span as much of the space of valid outputs as possible.




Figure 2: Each decoding strategy is plotted, showing that human-perceived quality is negatively correlated with diversity. The Pearson Correlation coefficients between each statistic and the average of fluency, coherence, and interestingness are shown in parentheses.

Prompt: Look, nobody knows we did it.		
RS 0.5 I don't know what you're talking about. What's the matter with you? I don't know what it is. I don't think so. He's got to get out of here.	Standard BS We've got to get out of here. What do you mean? I don't think it's a good idea. I don't know what to say. I don't know what's going on.	NPAD0.3 BS I don't think it's a good idea. I don't know what to say. I don't know what's going on. I don't know what to do. I don't know what's going on here.
RS 1.0 I can't find it. They're our ships. It's all right anyone is the right to interfere. We didn't have a plan I engineered a policy. Same time you pick us up at six and get we.	Standard BS with PDC I don't know! I don't think so. What do you mean? Why didn't you tell me? That's why we're here.	Cluster5 BS I don't know why. What do you mean? I don't think so. How do you know that? I'll tell you what.
RS 1.0,top10 I don't know what else to do. It doesn't have to be that way! We're in the air! I've seen a guy in his place in a it. And I'm not we any more.	RS 1.0,top10 with PDC What do you mean? I don't think so. That's why I'm here. It's all right we. We've been through this before.	Top3Cap BS We've got to get out of here. What do you mean? I don't think it's a good idea. I don't know what to say. I don't know what's going on.

Table 4: Responses to an example prompt for selected methods. More examples can be seen in the appendix.

However, we find there to be a marked trade-off between diversity and quality. This can be seen in Figure 2, where we plot the human-judged quality score for each dialog experiment against our primary diversity descriptive statistics. Fluency and adequacy are both strongly negatively correlated with diversity. While we had expected interestingness to be positively correlated with diversity, the fact that it is not suggests that existing diversity statistics are insufficient for capturing what it means to humans for outcomes to be interesting.

Likewise, in image captioning, the mean SPICE score of the 10 candidate captions (averaged over all examples for each experimental setting) is strongly anti-correlated with diversity, with a Pearson correlation coefficient of -0.83 with the Ent-4

measure and -0.84 with Dist-2. Clearly it remains an open challenge to generate a diverse set of image captions that are all high-quality.

When researchers choose to use a diverse decoding strategy, they must decide where on the quality-diversity tradeoff they would like to lie; selecting an optimal method depends strongly on one's tolerance for errors. In machine translation, where mistakes could severely impact coherence, beam search-based methods, which tend to result in better fluency and coherence, but worse diversity might be preferred. In more open-ended applications, where novel text is of greater importance, increased diversity could be worth the fluency and coherence hit. As state-of-the-art models continue to improve, one would hope that the quality cost of

encouraging diversity will continue to decrease.

In the interest of reporting a single overall best method for each task, we computed a sum-of-ranks score for each method. For dialog, we ranked the methods each by fluency, coherence, interestingness, and Ent-4, and then took a weighted sum of the four ranks, with 50% of the weight assigned to Ent-4, and 50% distributed evenly among the human evaluation ranks. Overall, clustered beam search and standard BS (beam size 100, PDC) have the best scores, followed by clustered beam search (beam size 10). Similarly, for image captioning, we rank the methods by their mean SPICE score and by Ent-4. Summing these ranks, random sampling (temp 1.0, top-10 capping, PDC) came in first. Standard beam search, Hamming Diversity beam search, and Top- g capping beam search (beam size 10) tied for second.

6.2 Random Sampling-based Methods

Higher sampling temperatures result in both an increase in diversity in generated responses and a reduction in overall quality. In the dialog domain, evaluators consistently rate the responses sampled with temperature 1.0 to have worse fluency, coherence, and interestingness when those sampled with temperature 0.5. In the image captioning domain, lower temperature improves automatic evaluation metrics for quality while reducing diversity.

For dialog, restricting sampling to the top-10 vocabulary words is a more effective strategy than adjusting temperature for ensuring balance between the quality and diversity of outputs. Top-10 random sampling has the highest fluency, coherence, and interestingness, as well as significantly lower perplexity than other random sampling methods. However, this trend did not extend to image captioning, where top-10 random sampling results in both worse SPICE scores and lower diversity measures than setting the temperature to 0.7. This may be because image captioning is a less ambiguous task than open-ended dialog, leading to a better-trained model that puts more probability mass on high-quality vocabulary words, ameliorating the challenge top- c filtering is designed to eliminate: that of a long tail of low probability vocabulary words taking up a large amount of probability mass.

6.3 Beam Search-based Methods

For dialog, clustered beam search (Cluster5 BS) performs the best of all beam search methods in

terms of human-judged interestingness. It ties for best with NPAD0.3BS on fluency and ties with Standard BS on coherence. Iterative beam search (Iter5 BS) achieves the greatest diversity, but at the expensive of quality. It has the lowest human-judged coherence among beam search methods; thus, we do not evaluate this method on image captioning. For image captioning, Cluster5 BS has the highest diversity among beam search methods, but this difference is quite small. Cluster5 BS also has the highest SPICE@10 score, indicating it is the best method for generating at least one high quality candidate. However, Top3Cap BS results in the highest mean SPICE score, suggesting it is best at ensuring all outputs are reasonable quality.

6.4 Effect of Over-sampling

In our experiments, we explore over-sampling 100 outputs, and then either using post-decoding clustering (PDC) or re-ranking by log-likelihood to filter these 100 down to 10 diverse outputs.

In the dialog domain, this over-sampling approach is a definite win. When over-sampling with random sampling both methods of filtering substantially improve human judgements of fluency and adequacy compared to random sampling only 10 outputs. However, interestingness scores go down, and while the outputs are still more diverse than beam search-based methods, they are less diverse than random sampling without filtering. In the beam search methods that use a beam size of 100 then filter down to 10, human-judged quality is on par with beam size 10 results, but diversity is considerably higher.

When comparing the two types of filtering, PDC results in higher interestingness and diversity statistics, while log-likelihood re-ranking improves fluency and adequacy. This again demonstrates the trade-off between quality and diversity.⁷

For image captioning, over-sampling with reranking does not consistently improve quality as it does in the dialog domain. Mean SPICE score is improved for random sampling but not for beam search. SPICE@1 becomes worse for both random sampling and decoding, while SPICE@10 improves for random sampling, and for beam search when PDC is applied. From these results,

⁷In the appendix, we show results with every method where we generate 10 samples; generate 100 samples followed by selecting the 10 most likely outputs; and generate 100 samples followed by post-decoding clustering to select 10 outputs.

we can conclude that over-sampling then ranking does not have a sizeable effect, either negative or positive, on quality. Moreover, the diversity of the captions generated by random sampling actually decreases when oversampling. The diversity of beam search-generated captions does improve with over-sampling.

While oversampling does generally improve outcomes on the diversity/quality tradeoff, it is more computationally expensive, particularly with beam search. Running PDC also requires generating sentence embeddings for every output, which adds additional computation time.

7 Additional Related Work

In this paper, we have compared a variety of post-training diversity-promoting algorithms. Here, we discuss other related works that instead promote diversity at train-time, as well as alternative quality evaluation methods. We also note that concurrent work has proposed nucleus sampling as an improvement to the sampling strategies discussed in this paper (Holtzman et al., 2019).

Diversity Promotion During Training Several works have attempted to encourage diversity during training by replacing the standard log-likelihood loss with a diversity-promoting objective. Li et al. (2016a) introduces an objective that maximizes mutual information between the source and target. Zhang et al. (2018) uses an adversarial information maximization approach to encourage generated text to be simultaneously informative and diverse. Xu et al. (2018) also uses an adversarial loss; their loss function rewards fluent text and penalizes repetitive text. We do not evaluate on these methods as they tend to be task-specific and difficult to implement. All of the diversity strategies we evaluate share the trait that they are agnostic to model architecture and to the data type of the input, as long as the output of the model is a probability distribution over tokens in a sequence.

Automatic Quality Evaluation An important part of this work is how to accurately measure not only the effect these methods have on candidate diversity, but also on the overall quality of the candidates. In choosing to report human scores and perplexity for the dialog domain, and SPICE for image captioning, we omitted some quality measures used in other papers.

For image captioning, BLEU (Papineni et al.,

2001), ROUGE (Lin, 2004), METEOR (Elliott and Keller, 2013), and CIDEr (Vedantam et al., 2015) scores are often reported, but SPICE has been shown to have higher correlation with human judgments (Anderson et al., 2016). In the dialog domain, single-reference BLEU score (Papineni et al., 2001) is sometimes used to measure response quality, but it has been shown to have little correlation with human-judged quality (Liu et al., 2016). Therefore, most works in dialog systems use human evaluation as the ultimate measure of quality (Li et al., 2016a; Sedoc et al., 2018)

8 Conclusion

In this work, we perform an analysis of post-training decoding strategies that attempt to promote diversity in conditional language models. We show how over-sampling outputs then filtering down to the desired number is an easy way to increase diversity. Due to the computational expense of running large beam searches, we recommend using random-sampling to over-sample. The relative effectiveness of the various decoding strategies differs for the two tasks we considered, which suggests that choice of optimal diverse decoding strategy is both task-specific and dependent on one’s tolerance for lower quality outputs.

While we have focused on evaluating each decoding strategy under the specifics reported to be the best in the original, further work is necessary to conclude whether observed differences in quality and diversity may simply be due to each work’s chosen hyperparameters. The ability to effectively generate a diverse set of responses while not degrading quality is extremely important in a variety of generation tasks, and is a crucial component to harnessing the power of state-of-the-art generative models.

9 Acknowledgements

We thank our anonymous reviewers for helpful feedback. We also thank Yun William Yu for assistance with statistical testing and proofreading. This material is based in part on research sponsored by DARPA under grant number HR0011-15-C-0115 (LORELEI). The U.S. Government is authorized to reproduce and distribute reprints for Governmental purposes. The views and conclusions in this publication are those of the authors and should not be seen as representing official endorsements of DARPA and the U.S. Government.

References

- Peter Anderson, Basura Fernando, Mark Johnson, and Stephen Gould. 2016. [Spice: Semantic propositional image caption evaluation](#). In *European Conference on Computer Vision*.
- Peter Anderson, Xiaodong He, Chris Buehler, Damien Teney, Mark Johnson, Stephen Gould, and Lei Zhang. 2018. [Bottom-up and top-down attention for image captioning and visual question answering](#). In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 6077–6086.
- Ashutosh Baheti, Alan Ritter, Jiwei Li, and Bill Dolan. 2018. Generating more interesting responses in neural conversation models with distributional constraints. In *Conference on Empirical Methods in Natural Language Processing (EMNLP 2018)*.
- Kyunghyun Cho. 2016. Noisy parallel approximate decoding for conditional recurrent language model.
- Sajal Choudhary, Perna Srivastava, Lyle H. Ungar, and João Sedoc. 2017. Domain aware neural dialog system. volume abs/1708.00897.
- Elizabeth Clark, Anne Spencer Ross, Chenhao Tan, Yangfeng Ji, and Noah A. Smith. 2018. [Creative writing with a machine in the loop: Case studies on slogans and stories](#). In *23rd International Conference on Intelligent User Interfaces, IUI '18*, pages 329–340, New York, NY, USA. ACM.
- Cristian Danescu-Niculescu-Mizil and Lillian Lee. 2011. [Chameleons in imagined conversations: A new approach to understanding coordination of linguistic style in dialogs](#). In *Proceedings of the 2nd Workshop on Cognitive Modeling and Computational Linguistics*, pages 76–87, Portland, Oregon, USA. Association for Computational Linguistics.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, , and Kristina Toutanova. 2018. Bert: Pre-training of deep bidirectional transformers for language understanding.
- Desmond Elliott and Frank Keller. 2013. [Image description using visual dependency representations](#). In *Proceedings of the 2013 Conference on Empirical Methods in Natural Language Processing*, pages 1292–1302, Seattle, Washington, USA. Association for Computational Linguistics.
- Angela Fan, Mike Lewis, and Yann Dauphin. 2018. [Hierarchical neural story generation](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 889–898, Melbourne, Australia. Association for Computational Linguistics.
- Kevin Gimpel, Dhruv Batra, Chris Dyer, and Gregory Shakhnarovich. 2013. [A systematic exploration of diversity in machine translation](#). In *Proceedings of the 2013 Conference on Empirical Methods in Natural Language Processing*, pages 1100–1111.
- Jiatao Gu, Kyunghyun Cho, and Victor O.K. Li. 2017. [Trainable greedy decoding for neural machine translation](#). In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 1968–1978, Copenhagen, Denmark. Association for Computational Linguistics.
- Tatsunori B. Hashimoto, Hugh Zhang, and Percy Liang. 2019. [Unifying human and statistical evaluation for natural language generation](#). *CoRR*, abs/1904.02792.
- Ari Holtzman, Jan Buys, Maxwell Forbes, and Yejin Choi. 2019. The curious case of neural text degeneration. *CoRR*, abs/1904.09751.
- Guillaume Klein, Yoon Kim, Yuntian Deng, Jean Senellart, and Alexander M. Rush. 2017. [OpenNMT: Open-source toolkit for neural machine translation](#). In *Proc. ACL*.
- Jonathan Krause, Justin Johnson, Ranjay Krishna, and Li Fei-Fei. 2017. A hierarchical approach for generating descriptive image paragraphs. In *Computer Vision and Pattern Recognition (CVPR), 2017 IEEE Conference on*, pages 3337–3345. IEEE.
- Reno Kriz, João Sedoc, Marianna Apidianaki, Carolina Zheng, Gaurav Kumar, Eleni Miltsakaki, and Chris Callison-Burch. 2019. Complexity-weighted loss and diverse reranking for sentence simplification.
- Ilya Kulikov, Alexander H Miller, Kyunghyun Cho, and Jason Weston. 2018. Importance of a search strategy in neural dialogue modelling.
- Quoc Le and Tomas Mikolov. 2014. Distributed representations of sentences and documents. In *Proceedings of the 31st International Conference on International Conference on Machine Learning - Volume 32, ICML'14*, pages 1188–1196.
- Jiwei Li, Michel Galley, Chris Brockett, Jianfeng Gao, and Bill Dolan. 2016a. [A diversity-promoting objective function for neural conversation models](#). In *Proceedings of the 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 110–119, San Diego, California. Association for Computational Linguistics.
- Jiwei Li and Dan Jurafsky. 2016. Mutual information and diverse decoding improve neural machine translation.
- Jiwei Li, Will Monroe, and Dan Jurafsky. 2016b. A simple, fast diverse decoding algorithm for neural generation.
- Chin-Yew Lin. 2004. [ROUGE: A package for automatic evaluation of summaries](#). In *Text Summarization Branches Out: Proceedings of the ACL-04 Workshop*, pages 74–81, Barcelona, Spain. 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 *European conference on computer vision*, pages 740–755. Springer.
- Chia-Wei Liu, Ryan Lowe, Iulian Serban, Mike Noseworthy, Laurent Charlin, and Joelle Pineau. 2016. [How NOT to evaluate your dialogue system: An empirical study of unsupervised evaluation metrics for dialogue response generation](#). In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 2122–2132, Austin, Texas. Association for Computational Linguistics.
- Ruotian Luo. 2017. An image captioning codebase in pytorch. <https://github.com/ruotianluo/ImageCaptioning.pytorch>.
- Thang Luong, Hieu Pham, and Christopher D. Manning. 2015. [Effective Approaches to Attention-based Neural Machine Translation](#). In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*, pages 1412–1421, Lisbon, Portugal.
- Ramesh Nallapati, Bowen Zhou, Cicero Nogueira dos Santos, aglar Gülehre, and Bing Xiang. 2016. [Abstractive Text Summarization using Sequence-to-sequence RNNs and Beyond](#). In *Proceedings of The 20th SIGNLL Conference on Computational Natural Language Learning (CoNLL)*, pages 280–290.
- Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. 2001. [Bleu: a method for automatic evaluation of machine translation](#). In *Association for Computational Linguistics*.
- Jeffrey Pennington, Richard Socher, and Christopher D. Manning. 2014. [Glove: Global vectors for word representation](#). In *Empirical Methods in Natural Language Processing (EMNLP)*, pages 1532–1543.
- Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, and Ilya Sutskever. 2019. Language models are unsupervised multitask learners. *OpenAI Blog*, 1:8.
- João Sedoc, Daphne Ippolito, Arun Kirubarajan, Jai Thirani, Lyle Ungar, and Chris Callison-Burch. 2018. [Chateval: A tool for the systematic evaluation of chatbots](#). In *Proceedings of the Workshop on Intelligent Interactive Systems and Language Generation (2IS&NLG)*, pages 42–44. Association for Computational Linguistics.
- Ilya Sutskever, Oriol Vinyals, and Quoc V Le. 2014. [Sequence to sequence learning with neural networks](#). In Z. Ghahramani, M. Welling, C. Cortes, N. D. Lawrence, and K. Q. Weinberger, editors, *Advances in Neural Information Processing Systems 27*, pages 3104–3112. Curran Associates, Inc.
- Yik-Cheung Tam, Jiachen Ding, Cheng Niu, and Jie Zhou. 2019. [Cluster-based beam search for pointer-generator chatbot grounded by knowledge](#). In *Dialog System Technology Challenges 7 at AAI 2019*.
- Jörg Tiedemann. 2009. News from opus-a collection of multilingual parallel corpora with tools and interfaces. In *Recent advances in natural language processing*, volume 5, pages 237–248.
- Ramakrishna Vedantam, C. Lawrence Zitnick, and Devi Parikh. 2015. [Cider: Consensus-based image description evaluation](#). pages 4566–4575.
- Ashwin K Vijayakumar, Michael Cogswell, Ramprasad R Selvaraju, Qing Sun, Stefan Lee, David Crandall, and Dhruv Batra. 2016. [Diverse beam search: Decoding diverse solutions from neural sequence models](#).
- Oriol Vinyals and Quoc V. Le. 2015. [A neural conversational model](#). volume abs/1506.05869.
- Jingjing Xu, Xuancheng Ren, Junyang Lin, and Xu Sun. 2018. [Diversity-promoting gan: A cross-entropy based generative adversarial network for diversified text generation](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 3940–3949.
- Yizhe Zhang, Michel Galley, Jianfeng Gao, Zhe Gan, Xiujun Li, Chris Brockett, and Bill Dolan. 2018. [Generating informative and diverse conversational responses via adversarial information maximization](#). In *Advances in Neural Information Processing Systems*, pages 1815–1825.