AVGOUT: A Simple Output-Probability Measure to Eliminate Dull Responses

Tong Niu, Mohit Bansal

UNC Chapel Hill

{tongn, mbansal}@cs.unc.edu

Abstract

Many sequence-to-sequence dialogue models tend to generate safe, uninformative responses. There have been various useful efforts on trying to eliminate them. However, these approaches either improve decoding algorithms during inference, rely on hand-crafted features, or employ complex models. In our work, we build dialogue models that are dynamically aware of what utterances or tokens are dull without any feature-engineering. Specifically, we start with a simple yet effective automatic metric, AVGOUT, which calculates the average output probability distribution of all time steps on the decoder side during training. This metric directly estimates which tokens are more likely to be generated, thus making it a faithful evaluation of the model diversity (i.e., for diverse models, the token probabilities should be more evenly distributed rather than peaked at a few dull tokens). We then leverage this novel metric to propose three models that promote diversity without losing relevance. The first model, MINAVGOUT, directly maximizes the diversity score through the output distributions of each batch; the second model, Label Fine-Tuning (LFT), prepends to the source sequence a label continuously scaled by the diversity score to control the diversity level; the third model, RL, adopts Reinforcement Learning and treats the diversity score as a reward signal. Moreover, we experiment with a hybrid model by combining the loss terms of MINAVGOUT and RL. All four models outperform their base LSTM-RNN model on both diversity and relevance by a large margin, and are comparable to or better than competitive baselines (also verified via human evaluation). Moreover, our approaches are orthogonal to the base model, making them applicable as an add-on to other emerging better dialogue models in the future.

1 Introduction

Many modern dialogue generation models use a sequence-to-sequence architecture as their backbone (Sordoni et al. 2015), following its success when applied to Machine Translation (MT) (Bahdanau, Cho, and Bengio 2015). However, dialogue tasks also have a requirement different from that of MT: the response not only has to be “correct” (coherent and relevant), but also needs to be diverse and informative. However, seq2seq has been reported by many previous

works to have low corpus-level diversity (Li et al. 2016a; Serban et al. 2016; Sordoni et al. 2015; Vinyals and Le 2015), as it tends to generate safe, terse, and uninformative responses, such as “*I don’t know*.” These responses unnecessarily make a dialogue system much less interactive than it should be.

To increase the diversity of dialogue responses, the first step is to faithfully evaluate how diverse a response is. There are metrics used by previous work that are correlated to diversity, but not strongly, such as ratio of distinct tokens (Li et al. 2016a) and response length (Baheti et al. 2018). However, a response can be long but extremely boring in meaning, such as “*I am sure that I don’t know about it*.”, or short but interesting (i.e., contains a lot of information), such as “*Dad was mean*.”. Only investigating discrete token output by the model is also not ideal, because these tokens are only a single realization of the model’s output probability distribution at each time step, which unavoidably loses valuable information indicated by the whole distribution. Li et al. (2016b) manually collect a shortlist of dull responses, and during training discourage the model from producing such utterances. However, an important drawback of hand-crafted rules is that the set of dull tokens or utterances is static, while in fact it usually evolves during training: when the current dull tokens are eliminated, another set of them might reveal themselves.

In our work,¹ we begin with a simple yet effective approach to measure how diverse a response is. This metric, which we name “Average Output Probability Distribution”, or AVGOUT, draws information directly from the training-in-session model itself. We calculate it by keeping track of the exponential average of all output probability distributions on the decoder side during training. This metric dynamically measures which tokens the model is biased toward without any hand-crafted rules, thus making it a faithful evaluation of the model diversity (i.e., for diverse models, the token probabilities should be more evenly distributed rather than peaked at a few dull tokens). In addition, since AVGOUT is a *one-dimensional categorical distribution* rather than a dimensionless numerical value like entropy, it naturally carries and conveys more information

Copyright © 2020, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

¹We will release all our code and model outputs.

about model diversity.

We then propose three models that leverage our novel metric to promote diversity in dialogue generation. The first MINAVGOUT model minimizes the dot product of current batch AVGOUT and the exponential average AVGOUT across batches, which encourages low-frequency tokens to be generated. The second LFT model uses a labeled transduction method and scales a “diversity label” by the diversity score of the ground-truth target sequence during training, while during testing can generate responses of different levels of diversity by tweaking the intended diversity score. The third RL model leverages reinforcement learning, where our novel metric is applied to discrete tokens and serve as a reward signal. In addition, since MINAVGOUT regularizes directly on the continuous distribution while RL calculates its reward based on discrete sampled tokens, we simply add up the loss terms of the two models, creating an even stronger hybrid model.

We first employ diverse automatic metrics, including Distinct-1 and -2 from previous work (Li et al. 2016a) and our novel metric DIVERISTY-*i*AUC (which calculates one minus the sum of normalized frequencies of the most frequent tokens produced by the model), plus activity/entity F1s, to evaluate the diversity and relevance of the generated responses. We then conduct human evaluations to verify that these models not only outperform their base model LSTM by a large margin, but are also comparable to or better than an advanced decoding algorithm MMI (Li et al. 2016a) and a very competitive model VHRED (Serban et al. 2017b) on the Ubuntu dataset.

2 AVGOUT as an Effective Diversity Metric

By only keeping a static shortlist of boring responses or tokens, one basically assumes that we humans should decide which tokens are dull. However, we argue that we should instead look from the model’s perspective to identify dull tokens, because even if the model outputs a word that we consider rare, including it in too many responses is still considered a dull behavior. Motivated by this thought experiment, we propose a novel metric, Average Output Probability Distribution (AVGOUT), that dynamically keeps track of which tokens the model is biased toward. To calculate this, during training, we average out all the output probability distributions for each time step of the decoder for the whole mini-batch. The resulting vector D' will reflect each token’s probability of being generated from the model’s perspective. Note that we do not use discrete ground-truth tokens to evaluate the model’s bias, because there is a fine distinction between the two: a statistics of frequency on ground-truth tokens is an evaluation of *the corpus’s bias*, while AVGOUT is an evaluation of *what bias the model has learned* because by generating dull responses more frequently than the training corpus has, it is the model itself that we should adjust. Also note that the reason we take the average is that a single output distribution will largely depend on the context and the previous target tokens (which are fed as inputs to the decoder during training), but on average the distribution should be a faithful evaluation on which words are more likely to be generated from model’s perspective.

To avoid batches that have AVGOUT significantly different from those of other batches, which would lead the model astray, we keep the exponential average of this metric across batches to make it less biased toward any specific batch. Let it be D . After training on a mini-batch and obtain D' , we update D like the following:

$$D \leftarrow \gamma D' + (1 - \gamma)D$$

where γ is 0.01 in our experiments.

Another consideration of AVGOUT is that theoretically we can have two choices. The first is to use the output distributions when we are teacher-forcing (i.e., only feeding ground-truth tokens); the other is to let the model use its own predictions during greedy/beam-search decoding or sampling. We reason that the former is a much better estimation of the model’s bias, because the latter will result in a cascading enlargement of the model bias due to the autoregressive nature of LSTM-RNN models (i.e., the tokens fed to the decoder are themselves also polluted by the model’s bias). Our early experimental results also agreed with the above reasoning.

Although we try to come up with the most faithful evaluation of how diverse a response is, our approach certainly has its drawbacks too. For example, using very frequent words but less frequent combinations of them may result in a good response which will be penalized by our metric. A natural solution to this is to also use bigram and trigram diversities and take a linear combination of them, which on a high-level is similar to BLEU (Papineni et al. 2002). However, considering even bigram distribution takes up $O(|V|^2)$ space and calculation time, hence we did not try it due to limited resources. However, as will be presented in Section 5, regularizing unigram distributions can already greatly help on higher-gram diversities, while also improving relevance.

3 Three Models to Leverage AVGOUT

AVGOUT can play at least three roles. First, it can be used to directly supervise output distribution during training; second, it can be used as a prior in labeled sequence transduction methods to control diversity of the generated response; and third, it can be used as a reward signal for Reinforcement Learning to encourage diverse sampled responses.² In this section, we begin with a base vanilla seq2seq model, and next present our three models to diversify responses based on AVGOUT.

Our base model LSTM is identical to that proposed by Bahdanau, Cho, and Bengio (2015), which consists of a single-layer bi-directional LSTM-RNN (Hochreiter and Schmidhuber 1997) encoder and a single-layer LSTM-RNN decoder with additive attention.

3.1 Regularization by Minimizing CONTINUOUS-AVGOUT

Our MINAVGOUT model (Figure 1) directly integrates AVGOUT into the loss function by summarizing it into a single numerical value named CONTINUOUS-AVGOUT. We do

²It can also be used to re-rank beams during inference, but we are much more interested in making the underlying model itself more diverse.

Figure 1: MinAvgOut model: use the dot product of average output distribution of the exponential average and the current batch to evaluate how diverse the current batch is.

this by taking the dot-product of D and D' (Figure 2). The intuition behind this simple calculation is that D can also be viewed as a set of weights which add up to 1.0, since it is a probability vector. By taking the dot product, we are actually calculating a weighted average of each probability in D' . To evaluate how diverse the model currently is, the duller tokens should obviously carry higher weights since they contribute more to the “dullness” of the whole utterance.³ Assuming that D is a column vector, the continuous diversity score is B_c , and the resulting extra loss term is L_B , the total loss L is given by:

$$B_c = 1 - D^T D'; L_B = (-1) * \alpha B_c; L = L_{ML} + L_B$$

where α is a coefficient to balance the regularization loss with the maximum likelihood loss (a.k.a. teacher forcing loss) L_{ML} . This is important because the regularization term continues to discourage the model from generating the ground-truth token, which we need to balance by ML loss to reduce the impact (otherwise the model will be led astray). Note that since D is a moving average which does not depend on the model parameters of the current mini-batch, only D' will result in gradient flow during back-propagation, which is what we intend.

3.2 Label-Fine-Tuning Model

We also borrow the continuous version of the Label-Fine-Tuning (LFT) model from Niu and Bansal (2018b), which is an extension of the discrete labeled sequence transduction methods (Kikuchi et al. 2016). The LFT model leverages a continuous label to serve as a prior for generating the target sequence. This label corresponds to an embedding just like a normal token does, but can be scaled by a continuous value. This model is applicable to our case because the diversity score of a response can also be viewed as a style, ranging from 0.0 to 1.0. Specifically, we add to the vocabulary a diversity label and scale its embedding vector with the intended diversity score of the target sequence. During training, this score is obtained by evaluating the diversity of the ground-truth target sequence (see Figure 3); during test time, we instead feed the model a diversity label scaled by a

³This linear combination is crucial, because if we naively add up all probabilities in D' and take the average, the result will be a useless constant $1.0/|V|$, where $|V|$ is the vocabulary size. This naive approach certainly cannot capture the diversity of a model.

Figure 2: An example of AVGOUT applied to a single token, which readily generalizes to multiple tokens within a response. We calculate diversity score of a continuous distribution through dot product. We sum up the values in the last graph. Note that although “turf” (the fourth word from the right in all three sub-figures) has higher probability in the current batch D' , it still contributes less than the word “is” to the overall diversity measure when taking the dot product, due to its low probability in the exponential average distribution D (i.e., lower weights). All probabilities are for illustration purpose and do not correspond to distributions from our models.

score of our choice (i.e., when we want the model to generate a more diverse response, we scale the label’s embedding by a higher score, while to generate a duller response, we scale the embedding by a lower one).

3.3 Reward-Based Reinforcement Learning

We also explore a model (see Figure 4) which regularizes on the discrete token level, because merely monitoring output probability distribution may ignore certain bad styles such as repetition (e.g. “I don’t don’t know.”). We use DISCRETE-AVGOUT to calculate the continuous diversity score of a discrete sequence. Let $\{G_1, G_2, \dots, G_{N_G}\}$ be a sequence of N_G tokens sampled by the model during training. Then from D , we extract the probabilities $\{P_1, P_2, \dots, P_{N_G}\}$ corresponding to each generated token. The diversity score B_d on these discrete tokens will be:

$$B_d = 1 - (P_1 + P_2 + \dots + P_{N_G})/N_{unique}$$

Figure 3: LFT model: the diversity label is scaled by the diversity score of the ground-truth target during training.

where N_{unique} is the number of unique tokens in the sampled sequence (see Figure 5). Note that this division explicitly discourages the model from outputting repeated tokens, because when that happens, the nominator will stay the same, while the denominator will decrease, resulting in a lower diversity score. Also note that MINAVGOUT can be complementary to RL since calculating diversity scores based on discrete tokens unavoidably loses valuable information from the output distribution before argmax is taken. In Section 5, we show with both automatic and human evaluations that this combination indeed achieves the best results among our models. Following Paulus, Xiong, and Socher (2018), our loss function consists of two terms. The first term is the Maximum Likelihood loss (L_{ML}); the other is the Reinforcement Learning loss (L_{RL}). The total loss L is then:

$$L = L_{ML} + \beta L_{RL}$$

$$L_{ML} = - \sum_{t=1}^n \log p(y_t^* | y_1^*, \dots, y_{t-1}^*, x)$$

$$L_{RL} = - (R - R_b) \sum_{t=1}^n \log p(y_t^s | y_1^s, \dots, y_{t-1}^s, x)$$

where β is a hyperparameter indicating how much weight we want to assign to the RL part of the loss, x is the source sequence, $\{y_t^*\}$ are the ground truth tokens and $\{y_t^s\}$ are the sampled tokens. We use a policy gradient method (Sutton et al. 2000) to calculate the RL loss. Specifically, we sample a response for each context x , and assign to it a reward R , which is equal to B_d because we want to encourage the model to be diverse. We also use a baseline R_b that helps reduce variance during training (Ranzato et al. 2016). In our case this baseline is again the exponential average of all B_d in previous mini-batches.

4 Experimental Setup

4.1 Dataset and Task

We use the task-oriented Ubuntu Dialogue dataset (Lowe et al. 2015), because it not only has F1 metrics to evaluate the relevance of responses, but the dialogues in them are also open-ended to allow enough space for diversity. We also chose this dataset because previous work, e.g., HRED (Serban et al. 2016) and VHRED (Serban et al. 2017b) both used Ubuntu to showcase their diversity-promotion models. Due to the popularity of this dataset, we were able to reproduce

Figure 4: RL model: the diversity score of the sampled response is fed back to the model as reward signal.

Figure 5: Diversity score calculation for a discrete sequence: $B_d = 1.0 - (0.05 + 0.2 + 0.1 + 0.01)/4$.

almost all models on this same dataset and have a meaningful comparison on their effectiveness of eliminating dullness. As future work, we plan to apply our models to other datasets where diversity is desired.

4.2 Automatic Evaluation

To measure the relevance of the model responses, we follow Serban et al. (2017a) and evaluate on F1’s for both activities (technical verbs, e.g., “upload”, “install”) and entities (technical nouns, e.g., “root”, “internet”). The F1’s are computed by mapping the ground-truth and model responses to their corresponding activity-entity representations (Serban et al. 2017a), who considered F1 to be “particularly suited for the goal-oriented Ubuntu Dialogue Corpus”. We did not evaluate on BLEU score (Papineni et al. 2002) because (Liu et al. 2016) showed that BLEU does not correlate well with dialogue quality. Lowe et al. (2017) also made similar observations on BLEU. To evaluate diversity, we employ two evaluation metrics from previous work, namely DISTINCT-1 and DISTINCT-2 (Li et al. 2016a). These are the ratios between the number of unique tokens and all tokens for unigrams and bigrams, respectively. In addition, we propose a novel diversity graph and its corresponding metric, which we name DIVERSITY-32 and DIVERSITY-AUC, respectively. We gather statistics of sentence, unigram, bigram and trigram, and sort their normalized frequencies from highest to lowest. Observing that all four graphs follow long-tail distributions, we only keep the highest 32 fre-

Figure 6: Diversity-32 graphs of all models. Curves with lower AUC correspond to more diverse models.

quencies and plot them.⁴ We then calculate one minus the Area under Curve (DIVERSITY-AUC) for each graph, which draws a high-level picture of how diverse a model is.

4.3 Human Evaluation

Although we proposed the effective AvgOut metric, we did find that the model sometimes still cheats to gain higher automatic diversity score. For example, as can be seen in the selected output examples (Section 5), the model tends to generate words with typo since these are rarer tokens as compared to their correct counterparts. This is unavoidable for noisy datasets like Ubuntu. Thus, without human evaluation, we can never be sure if our models are good or they only look good because our metrics are exploited.⁵

We thus also conducted human studies on Amazon MTurk to evaluate the generated responses with pairwise comparison for dialogue quality. We compare our models with an advanced decoding algorithm MMI (Li et al. 2016a) and two models, namely LSTM (Sordani et al. 2015) and VHRED (Serban et al. 2017b), both with additive attention. To our best knowledge, LSTM and VHRED were the primary models with which F1’s were reported on the Ubuntu dataset. Following Baheti et al. (2018), we employ two criteria: Plausibility and Content Richness. The first criterion measures whether the response is plausible given the context, while the second gauges whether the response is diverse

⁴One can also pick any reasonable number around 32 without loss of generality. Also note that we do not take into account what specific token each frequency corresponds to, because for example, although “*plethora*” is an unusual word, a model that outputs it all the time is still boring. Thus the frequencies are more important than what the tokens actually are.

⁵This is also true for Ubuntu: the human studies and F-1 automatic evaluation are complementary to each other because MTurk annotators are better at judging how coherent and informative a response is, while F1s concentrate more on the technical terms (i.e., whether important activities or entities are present in the response).

	Activity F1	Entity F1
LSTM	1.18	0.87
LSTM (<i>attn</i>)	3.30	1.81
HRED	4.34	2.22
VHRED	4.63	2.53
MMI	5.17	3.11
VHRED (<i>attn</i>)	5.94	3.52
Reranking-RL	5.67	3.73
MinAvgOut	6.83	4.77
LFT	7.55	5.05
RL	5.80	3.62
MinAvgOut + RL	8.05	5.47

Table 1: Automatic Evaluation Activity/Entity F1 results for baselines and our 3 models (*attn* means “with attention”). LSTM, HRED and VHRED are reported in Serban et al. (2017a), VHRED (*attn*) and Reranking-RL in Niu and Bansal (2018a), and the rest are produced by our work. All our four models have statistically significantly higher F1 values ($p < 0.001$) against VHRED (*attn*) and MMI.

and informative. The utterances were randomly shuffled to anonymize model identity. We only allowed annotators located in the US-located with at least an approval rate of 98% and 10,000 approved HITs. We collected 100 annotations in total after rejecting those completed by people who assign exactly the same score to all model responses. Since we evaluated 7 models, we collected 700 annotations in total, which came from a diverse pool of annotators.

4.4 Training Details

For each of the three models, the hidden size of the encoder is 256, while the decoder hidden size is 512. For MINAVGOUT, the coefficient of the regularization loss term α is 100.0; For LFT, during inference we feed a score of 0.015 since it achieves a good balance between response coherence and diversity. For RL, the coefficient of the RL term β is 100.0. For the hybrid model MINAVGOUT + RL, α and β share a coefficient of 50.0.

5 Results and Analysis

5.1 Automatic Evaluation Results

We employ several complementary metrics to capture different aspects of the model. The F1 results are shown in Table 1.⁶ Among all single models, LFT performs the best, followed by MINAVGOUT. RL is also comparable with previous state-of-the-art models VHRED (*attn*) and RERANKING-RL. We think that this is because LFT exerts no force in pulling the model predictions away from the ground-truth tokens, but rather just makes itself aware of how dull each response is. Consequently, its responses appear more relevant than the other two approaches. Moreover, the hybrid model (last row) outperforms all other mod-

⁶Note that the F1 scores for this task are overall low because the conversations in the Ubuntu dataset are all open-ended. This is unlike tasks such as Question Answering where there is usually a correct response.

	<i>i</i> AUC-s	<i>i</i> AUC-1	<i>i</i> AUC-2	<i>i</i> AUC-3	<i>i</i> AUC-avg	Distinct-1	Distinct-2
LSTM (<i>attn</i>)	0.545	0.310	0.614	0.665	0.5335	0.0074	0.0298
MMI	0.794	0.320	0.645	0.746	0.6263	0.0076	0.0312
VHRED (<i>attn</i>)	0.832	0.364	0.725	0.836	0.6893	0.0076	0.0474
MinAvgOut	0.902	0.428	0.754	0.838	0.7305	0.0082	0.0507
LFT	0.928	0.328	0.693	0.812	0.6903	0.0056	0.0301
RL	0.848	0.328	0.662	0.769	0.6518	0.0074	0.0354
MinAvgOut+RL	0.916	0.396	0.736	0.832	0.7200	0.0074	0.0442

Table 2: Automatic Evaluation results for the baselines and our proposed models (“*i*AUC” means “inverted AUC”, or “1 - AUC”; “*attn*” means “with attention”; “s”, “1”, “2” and “3” correspond to “sentence-level”, “unigram”, “bigram” and “trigram”, respectively; *i*AUC-avg is the average of all the other AUC columns). Best results are boldfaced. We do not calculate *p*-value because it does not apply to corpus-level metrics.

	Plaus.	Rich	Avg.	ScIDiff.
LSTM (<i>attn</i>)	3.46	2.62	3.04	0.28
MMI	3.57	2.92	3.25	0.20
VHRED (<i>attn</i>)	3.54	3.12	3.33	0.13
MinAvgOut	3.62	3.34	3.48	0.08
LFT	3.83	3.47	3.65	0.10
RL	3.61	2.88	3.25	0.22
MinAvgOut+RL	3.67	3.23	3.45	0.13

Table 3: Human Evaluation results for all the models we produce on Plausibility, Richness, average of the two, and scaled difference (difference between them divided by their average). Best Results are boldfaced. Note that for the last column, lower is better since we want balance between Plausibility and Content Richness. All results are pair-wise statistically significantly different with $p < 0.05$, except between MINAVGOUT and RL on Plausibility, and between MINAVGOUT and MINAVGOUT+RL on Average.

els by a large margin. One might expect that minimizing AVGOUT causes the models to move further away from the ground-truth tokens, so that it will hurt relevance. However, our F1 results show that as the responses become more diverse, they are more likely to include information more related and specific to the input contexts, which actually makes the model gain on both diversity and relevance. This will be further confirmed by the output examples in Table 4.

We also present DIVERSITY-32 graphs (Figure 6) and report DIVERSITY-AUC as well as DISTINCT-1 and -2 for each model (Table 2). We can see that all our models have significantly better sentence-level diversity than VHRED, let alone LSTM. For unigram diversity, they are also better than LSTM, though hard to distinguish from VHRED. Both bigram and trigram graphs reveal that all models are more diverse than LSTM, except that RL shows lower diversity than the other models, which agree with our F1 results. Note that since our models are only trained based on unigram output distributions, the bigram and trigram diversities are still far away from that of the ground-truth, which points to future direction. That said, the table does show that encouraging unigram diversity can already have positive influence on higher grams as well. Also note that the hybrid model (last row) does not achieve the best result in terms

of diversity. We hypothesize that this is because RL, which is usually harder to optimize than ML losses, faces exacerbated issues when combined with a strong MINAVGOUT loss, which tries to pull the model output distribution away from the token distribution in the training corpus.

Neither DISTINCT-1 nor -2 correlates well with our observation and evaluation of diversity and relevance. We reason that this is because these metrics only capture how many distinct tokens are used rather than each token’s frequency, which is easier to be exploited because whether each token is used unnecessarily often (a strong sign of dullness) is not reflected in this measure.

5.2 Human Evaluation Results

As mentioned in experimental setup, we conducted human evaluations on our models for both Plausibility and Content Richness, as well as calculating their average (to show overall score) and their difference (to show balance between the two criteria) (Table 3). We can see from the table that all our models are statistically significantly better than the baseline models on both Plausibility and Content Richness, except that RL is slightly weaker on Content Richness, which agrees with the trend in automatic evaluations. Although MINAVGOUT+RL model only ranks the second on average score (statistically equivalent to MINAVGOUT) in human evaluation, it achieves a good balance, and it also ranks the second in automatic diversity and the first in F1 values. We thus consider it to be our best model.

5.3 Selected Output Examples

We present two selected examples of generated responses from the investigated models (Table 4). We can see that all our models learn to attend well to the context, generating coherent and informative responses.

6 Related Work

6.1 Measurements of Response Diversity

Multiple metrics and approaches have been proposed to measure dialogue diversity. Some focus more on how similar the responses are to the ground-truth sequences, such as Word Error Rate (Serban et al. 2016) and BLEU (Galley et al. 2015), while the others explicitly have diversity in mind when being created, such as DISTINCT-1 and -2 (Li et al.

Context-X	do you think almoxarife 's network d/l idea is a good next step ?
Context-Y	the error message is from partman so its related to the disk partitions
Seq2seq	what do you mean ?
VRHED (attn)	it 's not showing up in fdisk -l
MMI	it 's not a disk , it 's a windows machine
MinAvgOut	I have a bunch of stuff in there , but I don't know how to do that
LFT	it says " no such file or directory "
RL	thanks for the tip
MinAvgOut + RL	I did that , but it says that it is not mounted
Context-X	im using the ubuntu default one and depending on which I install it should set the home why
Context-Y	run sudo update-alternatives -- config java then ? that should set the JAVA_HOME
Seq2seq	env — grep -i java
VRHED (attn)	what do you mean ?
MMI	i dont know how to do that
MinAvgOut	thanks
LFT	it says java is already the newest version
RL	it says that the java version is not installed
MinAvgOut + RL	it says " java 1.4 . 2 "
	ok , so what is the output of : sudo update-java-alternatives -l — grep java

Table 4: Selected output examples from all models. Context-X and -Y are given as model inputs during inference.

2016a). The key difference between AvgOut and the previous work is that first, our metric is dynamic with no feature-engineering; second, ours is versatile enough to be applied to both continuous distributions and discrete sequences, while theirs are only for discrete tokens; third, ours can be used for both sentence-level and corpus-level evaluation, while theirs are only meaningful as corpus-level metrics because they measure the extent of repetition across responses rather than for a single response.

6.2 Diversity-Promoting Dialogue Models

Researchers have different opinions on why dull responses are generated, which lead to various solutions. They can be roughly divided into four categories. The first category considers using conditional likelihood as a decoding objective the culprit (Baheti et al. 2018; Li et al. 2016a; Li, Monroe, and Jurafsky 2017; Shao et al. 2017). They thus focus on improving the decoding algorithm during training. The second category traces the cause of the low-diversity problem back to the lack of model variability. They then adopt Variational Autoencoders and rely on sampling from a latent random variable as an additional prior to the decoder (Serban et al. 2017b; Zhao, Zhao, and Eskenazi 2017; Gao et al. 2019). The third category thinks that the issue is a lack of universal background knowledge and common sense beyond the input context. They consequently aim to integrate prior knowledge into the generation process (Raghu, Gupta, and others 2019; Liu et al. 2018; Pei and Li 2018; Kryściński et al. 2018). The fourth category believes that the underlying model itself needs improvement. Some use hierarchical LSTM-RNN to encourage the model to capture high-level context (Serban et al. 2016); some use more advanced attention mechanism such as multi-head attention (Tao et al. 2018); and some use either more complicated architectures or models prone to degeneracies, such as

Generative Adversarial Networks (Li et al. 2017), Deep Reinforcement Learning (Li et al. 2016b) and Mixture Models (Shen et al. 2019). Our RL model has the same architecture as the Reinforcement Learning model, except with different rewards. Jiang and de Rijke (2018) consider the reason for dull responses as the model's over-confidence. They then propose to add to the loss function a regularization term to maximize the entropy of the output probability distribution. Interestingly, they only proposed this simple approach rather than actually implementing it. Our MINAVGOUT approach is related to their idea. Our approach is also related to posterior regularization (Mann and McCallum 2008; Ganchev et al. 2010; Zhu, Chen, and Xing 2014), but our work is neural-based.

7 Conclusion

We proposed a novel measure AVGOUT to dynamically evaluate how diverse a model or a response is based on the models' own parameters, which themselves evolve during training. We then leveraged this effective measure to train three models, plus a hybrid model, to eliminate dull responses for dialogue generation tasks. In addition, we designed novel automatic metrics to evaluate the trained models on diversity, in addition to the ones from previous work. Both automatic and human evaluations consolidated that our models are able to generate more diverse and relevant responses, even when compared with state-of-the-art approaches. For future work, we plan to apply these models to different generative tasks where diversity is desired.

Acknowledgments

We thank the reviewers for their helpful comments. This work was supported by NSF-CAREER Award #1846185, ONR #N00014-18-1-2871, and awards from Google, Face-

book, Salesforce (views are not of the funding agency).

References

- Bahdanau, D.; Cho, K.; and Bengio, Y. 2015. Neural machine translation by jointly learning to align and translate. In *Proceedings of ICLR*.
- Baheti, A.; Ritter, A.; Li, J.; and Dolan, B. 2018. Generating more interesting responses in neural conversation models with distributional constraints. In *Proceedings of EMNLP*.
- Galley, M.; Brockett, C.; Sordani, A.; Ji, Y.; Auli, M.; Quirk, C.; Mitchell, M.; Gao, J.; and Dolan, B. 2015. deltableu: A discriminative metric for generation tasks with intrinsically diverse targets. In *Proceedings of ACL*.
- Ganchev, K.; Gillenwater, J.; Taskar, B.; et al. 2010. Posterior regularization for structured latent variable models. *JMLR* 11(Jul):2001–2049.
- Gao, X.; Lee, S.; Zhang, Y.; Brockett, C.; Galley, M.; Gao, J.; and Dolan, B. 2019. Jointly optimizing diversity and relevance in neural response generation. In *NAACL*.
- Hochreiter, S., and Schmidhuber, J. 1997. Long short-term memory. *Neural computation* 9(8):1735–1780.
- Jiang, S., and de Rijke, M. 2018. Why are sequence-to-sequence models so dull? understanding the low-diversity problem of chatbots. In *Proceedings of EMNLP*.
- Kikuchi, Y.; Neubig, G.; Sasano, R.; Takamura, H.; and Okumura, M. 2016. Controlling output length in neural encoder-decoders. In *EMNLP*, 1328–1338.
- Kryściński, W.; Paulus, R.; Xiong, C.; and Socher, R. 2018. Improving abstraction in text summarization. In *ACL*.
- Li, J.; Galley, M.; Brockett, C.; Gao, J.; and Dolan, B. 2016a. A diversity-promoting objective function for neural conversation models. In *Proceedings of NAACL*.
- Li, J.; Monroe, W.; Ritter, A.; Galley, M.; Gao, J.; and Jurafsky, D. 2016b. Deep reinforcement learning for dialogue generation. In *Proceedings of ACL*.
- Li, J.; Monroe, W.; Shi, T.; Jean, S.; Ritter, A.; and Jurafsky, D. 2017. Adversarial learning for neural dialogue generation. In *Proceedings of ACL*.
- Li, J.; Monroe, W.; and Jurafsky, D. 2017. Learning to decode for future success. *arXiv preprint arXiv:1701.06549*.
- Liu, C.-W.; Lowe, R.; Serban, I. V.; Noseworthy, M.; Charlin, L.; and Pineau, J. 2016. How NOT to evaluate your dialogue system: An empirical study of unsupervised evaluation metrics for dialogue response generation. In *EMNLP*.
- Liu, S.; Chen, H.; Ren, Z.; Feng, Y.; Liu, Q.; and Yin, D. 2018. Knowledge diffusion for neural dialogue generation. In *Proceedings of ACL*, volume 1, 1489–1498.
- Lowe, R.; Pow, N.; Serban, I.; and Pineau, J. 2015. The ubuntu dialogue corpus: A large dataset for research in unstructured multi-turn dialogue systems. In *SIGDIAL*.
- Lowe, R.; Noseworthy, M.; Serban, I. V.; Angelard-Gontier, N.; Bengio, Y.; and Pineau, J. 2017. Towards an automatic turing test: Learning to evaluate dialogue responses. In *ACL*.
- Mann, G. S., and McCallum, A. 2008. Generalized expectation criteria for semi-supervised learning of conditional random fields. *ACL-08: HLT* 870.
- Niu, T., and Bansal, M. 2018a. Adversarial over-sensitivity and over-stability strategies for dialogue models. In *CoNLL*.
- Niu, T., and Bansal, M. 2018b. Polite dialogue generation without parallel data. *TACL* 6:373–389.
- Papineni, K.; Roukos, S.; Ward, T.; and Zhu, W.-J. 2002. BLEU: A method for automatic evaluation of machine translation. In *Proceedings of ACL*, 311–318.
- Paulus, R.; Xiong, C.; and Socher, R. 2018. A deep reinforced model for abstractive summarization. In *ICLR*.
- Pei, J., and Li, C. 2018. S2spmn: a simple and effective framework for response generation with relevant information. In *Proceedings of EMNLP*, 745–750.
- Raghu, D.; Gupta, N.; et al. 2019. Hierarchical pointer memory network for task oriented dialogue. In *NAACL*.
- Ranzato, M.; Chopra, S.; Auli, M.; and Zaremba, W. 2016. Sequence level training with recurrent neural networks. In *Proceedings of ICLR*.
- Serban, I. V.; Sordani, A.; Bengio, Y.; Courville, A.; and Pineau, J. 2016. Building end-to-end dialogue systems using generative hierarchical neural network models. In *AAAI*.
- Serban, I. V.; Klinger, T.; Tesauro, G.; Talamadupula, K.; Zhou, B.; Bengio, Y.; and Courville, A. C. 2017a. Multiresolution recurrent neural networks: An application to dialogue response generation. In *Proceedings of AAAI*, 3288–3294.
- Serban, I. V.; Sordani, A.; Lowe, R.; Charlin, L.; Pineau, J.; Courville, A.; and Bengio, Y. 2017b. A hierarchical latent variable encoder-decoder model for generating dialogues. In *AAAI*.
- Shao, L.; Gouws, S.; Britz, D.; Goldie, A.; Strophe, B.; and Kurzweil, R. 2017. Generating high-quality and informative conversation responses with sequence-to-sequence models. In *Proceedings of EMNLP*.
- Shen, T.; Ott, M.; Auli, M.; and Ranzato, M. 2019. Mixture models for diverse machine translation: Tricks of the trade. In *Proceedings of ICML*.
- Sordani, A.; Galley, M.; Auli, M.; Brockett, C.; Ji, Y.; Mitchell, M.; Nie, J.-Y.; Gao, J.; and Dolan, B. 2015. A neural network approach to context-sensitive generation of conversational responses. In *Proceedings of NAACL*.
- Sutton, R. S.; McAllester, D.; Singh, S.; and Mansour, Y. 2000. Policy gradient methods for reinforcement learning with function approximation. In *NeurIPS*.
- Tao, C.; Gao, S.; Shang, M.; Wu, W.; Zhao, D.; and Yan, R. 2018. Get the point of my utterance! learning towards effective responses with multi-head attention mechanism. In *Proceedings of IJCAI*, 4418–4424.
- Vinyals, O., and Le, Q. 2015. A neural conversational model. In *Proceedings of ICML*.
- Zhao, T.; Zhao, R.; and Eskenazi, M. 2017. Learning discourse-level diversity for neural dialog models using conditional variational autoencoders. In *Proceedings of ACL*.
- Zhu, J.; Chen, N.; and Xing, E. P. 2014. Bayesian inference with posterior regularization and applications to infinite latent svms. *JMLR* 15(1):1799–1847.