

Diversifying Dialog Generation via Adaptive Label Smoothing

Yida Wang^{1,2*}, Yinhe Zheng^{1,3*}, Yong Jiang², Minlie Huang^{1 †}

¹ The CoAI group, DCST, Institute for Artificial Intelligence, State Key Lab of Intelligent Technology and Systems, Beijing National Research Center for Information Science and Technology, Tsinghua University, Beijing, China

² Tsinghua-Berkeley Shenzhen Institute, Tsinghua Shenzhen International Graduate School, Tsinghua University, Shenzhen, China

³ Samsung Research China - Beijing (SRC-B)

wangyd18@mails.tsinghua.edu.cn, yh.zheng@samsung.com,
jiangy@sz.tsinghua.edu.cn, aihuang@tsinghua.edu.cn

Abstract

Neural dialogue generation models trained with the one-hot target distribution suffer from the over-confidence issue, which leads to poor generation diversity as widely reported in the literature. Although existing approaches such as label smoothing can alleviate this issue, they fail to adapt to diverse dialog contexts. In this paper, we propose an **Adaptive Label Smoothing (AdaLabel)** approach that can adaptively estimate a target label distribution at each time step for different contexts. The maximum probability in the predicted distribution is used to modify the soft target distribution produced by a novel light-weight bi-directional decoder module. The resulting target distribution is aware of both previous and future contexts and is adjusted to avoid over-training the dialogue model. Our model can be trained in an end-to-end manner. Extensive experiments on two benchmark datasets show that our approach outperforms various competitive baselines in producing diverse responses.

1 Introduction

The success of neural models has greatly advanced the research of dialog generation (Huang et al., 2020; Wang et al., 2020; Zhang et al., 2020). However, most of these models suffer from a low-diversity issue where models tend to generate bland and generic responses such as *I don't know* or *I'm OK* (Li et al., 2016). Although various approaches have been proposed to tackle this issue (Li et al., 2016; Zhao et al., 2017; Du et al., 2018; Zhou et al., 2018; Welleck et al., 2020; Zheng et al., 2020b), there are still remarkable gaps between responses generated by neural models and those from humans (Holtzman et al., 2020). Further, some existing methods may even harm the fluency or coherence when improving the diversity of generated

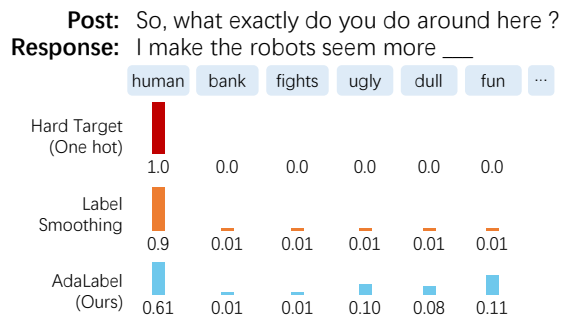


Figure 1: A dialogue sampled from the OpenSubtitles dataset. We demonstrate the hard target, label smoothing, and **Adaptive Label Smoothing** approach when learning to predict the next word (“human”).

responses. (Ippolito et al., 2019; Massarelli et al., 2020; Zheng et al., 2020a).

Recently, Jiang and de Rijke (2018); Jiang et al. (2019) show that there is a strong connection between the low-diversity problem and the over-confidence issue. i.e., over-confident dialogue models tend to produce low-diversity responses. One of the reasons can be attributed to the supervision target. Specifically, training a dialogue generation model with the Maximum Likelihood Estimation (MLE) objective under the hard target (i.e., one-hot distribution as ground truth) makes the model favor high-frequency tokens and produce over-confident probability estimation (Gowda and May, 2020), which ultimately leads to poor calibration (Mukhoti et al., 2020), and thus low diversity (Jiang et al., 2019). Hinton et al. (2015) and Yang et al. (2018) suggest that the ideal training target should be a soft target that assigns probability mass on multiple valid candidates (see Figure 1). With such a soft target, the over-confidence issue can be alleviated (Müller et al., 2019), and thus the diversity of the output responses can be improved.

Unfortunately, the ideal soft target is challenging to obtain. Early works try to tackle this issue

* Equal contribution

† Corresponding Author: aihuang@tsinghua.edu.cn

using label smoothing (Szegedy et al., 2016), i.e., a small probability is uniformly assigned to non-target words. However, the target distribution constructed in this way is far from ideal: **First**, the probability of the target word is chosen manually and fixed, which cannot adapt to different contexts. However, as Holtzman et al. (2020) demonstrated, human text distribution exhibits remarkable fluctuations in the per-token perplexity. We argue that different target probabilities should be used for different contexts. **Second**, the uniform assignment of the probability mass on non-target words ignores the semantic relationship between the context and each word. Ideally, a word should receive more probability mass if it is more relevant to the context. For the example shown in Figure 1, word “*fun*” is more likely to appear behind the context “*I make the robots seem more _*” than word “*bank*”.

To address the above issue, we propose an **Adaptive Label** smoothing (AdaLabel) method that can dynamically estimate a soft target distribution at each time step for different contexts. Specifically, for each target word y_t in the training data, the probability distribution predicted by the current model is first obtained. The maximum probability p_{max} in this distribution measures the confidence of the current prediction, i.e., a higher p_{max} means higher confidence for the current prediction. To avoid over-confidence, we use p_{max} as the supervision signal for the target word y_t in the training process so that the model will not be optimized towards y_t when it correctly predicts y_t . A word-level factor is also introduced to facilitate the learning of low-frequency words.

Moreover, we introduce a novel auxiliary decoder module D_a to produce the supervision signals for these non-target words in each training step. D_a only contains one transformer block, and it is optimized to predict words based on bi-directional contexts. A novel Target-Mask attention scheme is devised to prevent D_a from seeing the target word in the training process. This scheme also enables parallel training and inference of D_a .

We perform extensive experiments on two benchmark datasets: DailyDialog and OpenSubtitles. Our method outperforms various competitive baselines and significantly improves the diversity of generated responses while ensuring fluency and coherency. Our major contributions are summarized:

1. We propose AdaLabel, a method that can produce a soft target distribution considering the

current context and the model’s confidence. Specifically, AdaLabel ensures that the dialogue model will not be optimized toward the target word y_t if y_t has been correctly predicted. This prevents our model from being over-confident.

2. We introduce a light-weight bi-directional decoder that can produce context-aware supervision signals for non-target words. A novel Target-Mask attention scheme is devised to facilitate the parallel training and inference of this decoder.

3. Extensive experiments on two benchmark dialogue datasets with both automatic and human evaluation results show that our method helps to alleviate the model over-confident issue and significantly improves the model’s diversity.

2 Related work

Diversity Promotion: Existing approaches for solving the low diversity issue of neural dialogue models generally involve two categories:

The first category is training-based, where new training objectives are designed (Li et al., 2016; Zhang et al., 2018; Gao et al., 2019) or latent variables are introduced (Zhao et al., 2017; Zhou et al., 2018) in the dialogue model. Some methods also try to refine the training target used in the MLE loss (Choi et al., 2020; Jiang et al., 2019; Li et al., 2019), or directly penalize the trivial responses with auxiliary loss terms (Welleck et al., 2020; Li et al., 2020). Unlike these existing approaches, our method tries to adaptively adjust the training target by utilizing the current predictions.

The second category is decoding-based, in which different heuristic decoding rules are designed (Holtzman et al., 2020; Kulikov et al., 2019). Note that these decoding techniques are independent of the model setting, and our method can be used in combination with these techniques.

Confidence Calibration: Modern deep neural networks suffer from the over-confidence issue (Guo et al., 2017; Kumar and Sarawagi, 2019), and various remedies are proposed (Pereyra et al., 2017; Mukhoti et al., 2020; Lin et al., 2017). Following the work of Jiang and de Rijke (2018); Jiang et al. (2019), our method is proposed to tackle the over-confidence issue to improve the diversity of the generated responses. However, different from existing approaches, our method enables more flexible controls over the target distribution.

Knowledge Distillation: Another important technique similar to our work is knowledge distilla-

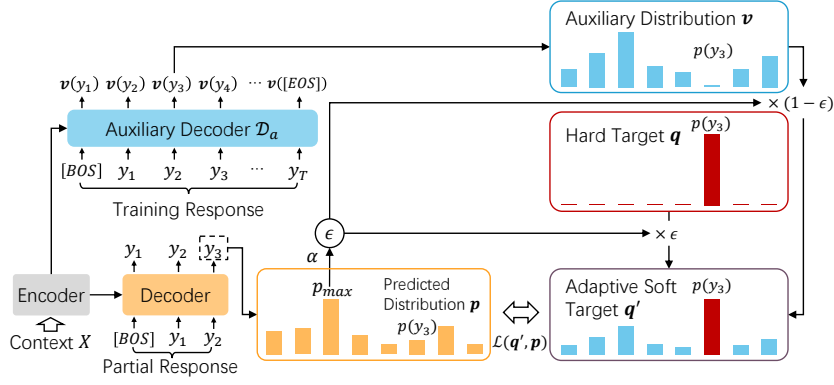


Figure 2: Overview of constructing the adaptive soft target \mathbf{q}' using AdaLabel: The maximum probability p_{max} in the predicted distribution \mathbf{p} is used to obtain an adaption factor ϵ , which is further used to combine the hard target \mathbf{q} and the auxiliary distribution \mathbf{v} to obtain \mathbf{q}' . A bi-directional auxiliary decoder \mathcal{D}_a is used to produce \mathbf{v} .

tion, in which a learned teacher model is distilled to a student model by minimizing a KL term (Hinton et al., 2015; Kim and Rush, 2016).

The most related work comparing to ours is the C-MLM approach (Chen et al., 2020), in which a BERT model is fine-tuned to be a teacher. Our approach and C-MLM’s primary difference is that our auxiliary decoder \mathcal{D}_a is a one layer module that is jointly trained with the dialogue model. However, the BERT teacher in C-MLM contains much more parameters, and it is trained using an expensive pre-trained and then fine-tuned process. Moreover, the target-masked attention scheme in \mathcal{D}_a enables parallel inferences of \mathbf{v} for each training sequence Y . In contrast, multiple independent forward passes are required for the BERT teacher.

3 Method

3.1 Background: MLE with Hard Target

The goal of generative dialogue modeling is to learn a conditional probability distribution $p(Y|X)$, where X is the dialogue context, $Y = y_1, \dots, y_T$ is a response word sequence, and $y_i \in \mathcal{V}$ is a word from the vocabulary \mathcal{V} . In an auto-regressive manner, $p(Y|X)$ is factorized as $\prod_t p(y_t|y_{<t}, X)$. For each target word y_t in the training sequence Y , a conventional MLE training approach try to optimize the following cross entropy loss:

$$\mathcal{L}(\mathbf{q}, \mathbf{p}) = - \sum_{w_k \in \mathcal{V}} q_k \log [p(w_k|y_{<t}, X)], \quad (1)$$

where \mathbf{q} is a one-hot distribution (i.e., a hard target) that assigns a probability of 1 for the target word y_t and 0 otherwise, i.e., $q_k = 1$ only when $w_k = y_t$. For simplicity of notation, we abbreviate

the dependency of y_t in the notation of each distribution in our paper, i.e., different target word y_t in Y corresponds to different values of \mathbf{q} and \mathbf{p} .

3.2 Method Overview

We propose to adaptively construct a soft target distribution \mathbf{q}' to replace \mathbf{q} in Eq. 1. Specifically,

$$\mathbf{q}' = \epsilon \cdot \mathbf{q} + (1 - \epsilon) \cdot \mathbf{v}, \quad (2)$$

where $\epsilon \in [0, 1]$ is an adaption factor, and \mathbf{v} is an auxiliary distribution vector that depends on the current time step. (see Figure 2 for an overview).

In this study, we constrain \mathbf{v} to assign zero probability for the target word y_t and non-zero probabilities for these non-target words $\mathcal{V}_{\neq y_t} = \{y_i | y_i \in \mathcal{V}, y_i \neq y_t\}$. This constraint allows us to explicitly control the supervisions assigned to y_t . Specifically, the first term $\epsilon \cdot \mathbf{q}$ and the second term $(1 - \epsilon) \cdot \mathbf{v}$ in Eq. 2 respectively determines how much probability \mathbf{q}' assigns to y_t and $\mathcal{V}_{\neq y_t}$. This setting differs from conventional knowledge distillation (Kim and Rush, 2016) because it facilitates more flexible controls over \mathbf{q}' , so that we can use the factor ϵ to determine the supervision signal provided for the target word y_t . The following sections detail how to compute ϵ and \mathbf{v} .

3.3 Target Word Probability

We control the probability of the target word y_t in \mathbf{p}' by manipulating the adaption factor ϵ in Eq. 2. Specifically, for a training dialogue pair $\langle X, Y \rangle$ and each target word $y_t \in Y$, the current distribution $p(\cdot|y_{<t}, X)$ is first calculated, and the maximum probability in this distribution is obtained:

$$p_{max} = \max_{w_k \in \mathcal{V}} p(w_k|y_{<t}, X). \quad (3)$$

ε is then obtained:

$$\varepsilon = \max(p_{max}, \lambda), \quad (4)$$

where λ serves as a lower-bound of ε (i.e., $\varepsilon \geq \lambda$).

The basic intuition behind Eq. 4 is to set $\varepsilon = p_{max}$ when p_{max} is reasonably large. This design prevents our model from receiving supervisions sharper than p_{max} , when the current prediction is confidence enough.

Further, to ensure that the target word y_t always receives the largest probability in \mathbf{q}' , i.e., to ensure $\varepsilon > (1 - \varepsilon) \cdot \max(\mathbf{v})$ (see Eq. 2), in which $\max(\mathbf{v})$ is the maximum probabilities for non-target words $\mathcal{V}_{\neq y_t}$, we have to enforce $\varepsilon > \frac{\max(\mathbf{v})}{1 + \max(\mathbf{v})}$. Thus we propose to calculate the lower-bound λ of ε as:

$$\lambda = \frac{\max(\mathbf{v})}{1 + \max(\mathbf{v})} + \eta, \quad (5)$$

where $\eta > 0$ is a hyper-parameter that controls the margin between the probability of the target word and non-target words in \mathbf{p}' .

To facilitate faster converge and better learning of low-probability words, an empirical factor $\alpha \in [0, 1]$ is further introduced to adjust the calculation of ε on the basis of Eq. 4:

$$\varepsilon = 1 - \alpha \cdot (1 - \max(p_{max}, \lambda)), \quad (6)$$

where α is calculated as the relative ratio to p_{max} :

$$\alpha = \left[\frac{p(y_t | y_{<t}, X)}{p_{max}} \right]^2, \quad (7)$$

where $p(y_t | y_{<t}, X)$ is the probability for the target word y_t . Note that Eq. 6 and Eq. 4 is equivalent if $\alpha = 1$. Intuitively, α accelerates the training of low-frequency words because if y_t is of low-frequency in the corpus, then y_t is usually under-trained and thus $p(y_t | y_{<t}, X)$ is generally small. This leads to a small α and thus increases the probability for y_t in \mathbf{p}' .

Note that ε , λ and α are all time-step specific variables, whereas η is a fixed hyper-parameter. This allows the values adapt to dynamic contexts. In our experiments, Eq. 6 is used to calculate ε .

3.4 Non-target Words Probabilities

The auxiliary distribution \mathbf{v} in Eq. 2 is calculated using an auxiliary decoder D_a , which is a single-layer transformer-based decoder that is jointly optimized with the generation model. Figure 3 shows the structure of D_a , in which a novel target-masked

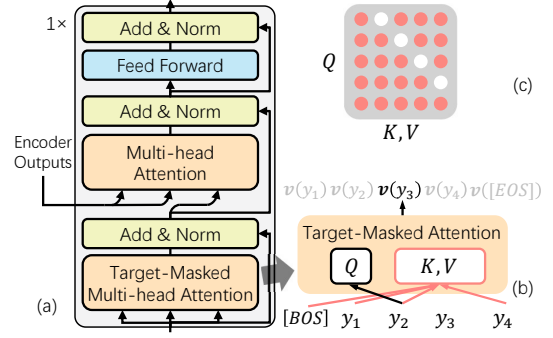


Figure 3: (a) The auxiliary decoder D_a ; (b) The target-masked attention scheme used to compute the auxiliary distribution \mathbf{v} for the target word y_3 , specifically, y_2 is used as the query and y_3 is masked; (c) The attention pattern used in the target-masked attention scheme, white dots represent masked positions.

attention scheme is devised to mask each target word y_t in the self attention module of the decoder when calculating the corresponding \mathbf{v} (see Figure 3b and 3c). In this way, bi-directional contexts can be utilized when predicting the auxiliary distribution \mathbf{v} for y_t . Moreover, it is important to use only one decoder layer in D_a because stacking multiple layers in D_a leaks the information of y_t to \mathbf{v} .

Note that using one layer in D_a does not necessarily downgrade its performance (Kasai et al., 2021). Our experiment results in Section 5.1 indicate that with the help of bi-directional contexts, the accuracy of D_a largely outperforms the unidirectional dialogue decoder that is much deeper than D_a . Moreover, for a training response Y , the structure of D_a enables us infer the auxiliary distribution in parallel for all the target words in Y within a single forward pass. This differs from the BERT teacher used by Chen et al. (2020), in which multiple independent forward passes are needed to get the teacher distributions for all the words in Y .

When training D_a , the following standard MLE loss is optimized for each target word y_t :

$$\mathcal{L}(\mathbf{q}, \mathbf{v}) = - \sum_{k=1}^{|\mathcal{V}|} q_k \log v_k, \quad (8)$$

in which the notation of q_k follows Eq. 1.

The outputs of D_a are used as the logits to infer \mathbf{v} to be further used in Eq. 2. Specifically, the logit of the target word y_t is masked to $-\infty$ before Softmax to ensure y_t always receives zero probability in \mathbf{v} . Moreover, we also follow the approach used by Tang et al. (2020) to truncate the head and tail of the remaining logits before inferring \mathbf{v} in Eq.

	Train	Valid	Test
DailyDialog	65.8K	6.13K	5.80K
OpenSubtitles	1.14M	20.0K	10.0K

Table 1: Dataset statistics.

2, i.e., all the logits are ranked in a descending order and only the logits ranked from n to m are kept while the rest logits are masked to $-\infty$. This masks the head and tail probabilities in v to zero. We argue that truncating the tail probabilities of v filters noises, and truncating the head probabilities of v encourages the dialogue model to focus more on low-probability words. In our experiments, we set $n = 2$ and $m = 500$. An extensive hyper-parameter search indicates that our method is not sensitive to the value of n and m .

There are two major differences between our auxiliary decoder D_a and the teacher model used in conventional knowledge distillation approaches: First, conventional teacher models usually carry more parameters than their students, whereas D_a is rather light-weight. Second, conventional teacher models are typically pre-trained before being utilized in the distillation process, whereas D_a is trained jointly with our dialogue model.

4 Experiments

4.1 Dataset

We use two benchmark datasets for open-domain dialogue generation: **DailyDialog** (Li et al., 2017) is a high-quality multi-turn dialogue dataset that is collected from daily conversations. **OpenSubtitles**¹ contains dialogues collected from movie subtitles. Moreover, we follow Li et al. (2016) and Jiang et al. (2019) to focus on short conversations, i.e., dialogues with posts or responses longer than 100 tokens are removed. See Table 1 for more details.

4.2 Implementation Details

The backbone of our model is the transformer-based sequence to sequence model (Vaswani et al., 2017), and most hyper-parameters follow Cai et al. (2020). Specifically, the encoder and decoder each contains 6 layers. Each layer has 8 attention heads, and the hidden size is set to 512. The auxiliary decoder D_a follows the same hyper-parameter setting as the dialogue decoder, but it only contains one layer. The WordPiece tokenizer provided by

¹<http://opus.nlpl.eu/OpenSubtitles.php>

BERT (Devlin et al., 2019) is used, and the Adam optimizer (Kingma and Ba, 2015) is employed to train our model from random initializations with a learning rate of $1e-4$. η in Eq. 5 is set to 0.2 for all datasets. See Appendix A for more details.²

4.3 Baselines

We compared our method with two groups of baselines that try to tackle the over-confidence issue.

The first group modifies the training target used to compute the loss function: **1) LS** (Szegedy et al., 2016): uses the label smoothing approach to construct a target distribution by adding the one-hot target and a uniform distribution; **2) FL** (Lin et al., 2017): uses the focal loss to down-weight well-classified tokens in each time step. **3) FACE** (Jiang et al., 2019): uses the frequency-aware cross-entropy loss to balance per-token training losses. Specifically, relative low losses are assigned to high-frequency words to explicitly tackle the over-confidence issue. We used the best performing “Pre-weigh” version in our experiments. **4) F²** (Choi et al., 2020): factorizes the target distribution based on the token frequencies.

The second group of baselines add some penalty term to the standard MLE loss: **5) CP** (Pereyra et al., 2017): a confidence penalty term is added to regularize the entropy of the model, so that over-confident predictions are penalized; **6) UL** (Welleck et al., 2020): an unlikelihood loss term is added to penalize the frequently generated words. **7) NL** (He and Glass, 2020): works similarly with baseline **UL** except a negative loss term is used instead of the unlikelihood loss term. **8) D2GPO** (Li et al., 2019): augments the MLE loss with a data-dependent gaussian prior objective to assign different losses for different non-target words.

We also compared to: **9) CE**: a vanilla Seq2Seq model trained with the cross-entropy loss. For fair comparisons, the C-MLM model proposed by Chen et al. (2020) is not used as our baseline since the BERT teacher in C-MLM requires a large amount of extra data to pre-train. Nevertheless, AdaLabel still surpasses C-MLM on various metrics (see Appendix F for more analysis).

All our baselines are adapted from the authors’ official codes with the same backbone architecture and hyper-parameters as our model (see details in Appendix B). Following the original setting, a train-

²Our code is available at: <https://github.com/lemon234071/AdaLabel>

Model	DailyDialog								OpenSubtitles							
	Dist-1, 2		Ent-1, 2		LF	BLEU-2,3,4			Dist-1, 2		Ent-1, 2		LF	BLEU-2,3,4		
CE	1.67	9.43	4.53	6.59	2.99	7.56	4.38	2.61	2.55	9.87	4.13	5.58	0.84	7.60	4.30	2.57
LS	1.48	8.78	4.48	6.55	2.44	7.98	4.68	2.86	2.77	13.08	4.45	6.57	0.51	8.91	5.57	3.84
FL	2.38	13.42	4.7	7.04	5.05	9.74	6.12	4.11	3.19	13.16	4.42	6.50	1.04	8.06	4.79	3.08
FACE	1.62	11.04	4.96	7.27	4.11	8.78	5.06	3.06	3.31	14.06	4.77	7.05	1.33	7.69	4.40	2.70
F ²	1.40	7.91	4.35	6.28	2.32	7.78	4.45	2.60	2.89	11.40	4.24	6.14	0.99	7.52	4.30	2.62
CP	2.35	12.91	4.64	6.89	4.07	9.06	5.68	3.79	3.11	12.72	4.36	6.35	0.98	8.06	4.82	3.12
UL	2.35	12.99	4.68	6.98	4.96	10.83	6.87	4.61	2.84	11.64	4.31	6.32	0.76	7.73	4.59	2.96
NL	1.66	9.18	4.47	6.58	4.30	9.83	5.83	3.60	3.24	12.98	4.42	6.49	1.08	7.56	4.38	2.71
D2GPo	1.26	8.06	4.43	6.48	2.20	8.30	4.82	2.93	2.07	11.01	4.32	6.36	0.19	8.41	5.08	3.35
AdaLabel	3.96	23.53	5.17	8.00	8.49	17.42	13.38	11.01	4.78	22.88	4.96	7.66	1.47	9.80	6.48	4.75
Human	6.59	37.74	5.67	8.91	13.7	N/A	N/A	N/A	8.62	43.16	5.89	9.36	4.75	N/A	N/A	N/A

Table 2: Automatic evaluation results (%). Best results among all the models are in bold.

and-refine strategy is used in baseline 3, 6, and 7, i.e., these baselines are refined based on **CE**. We follow the setting of [Jiang et al. \(2019\)](#) to use deterministic decoding scheme (particularly, greedy decoding) for our model and all baselines. Note that our method can be adapted to other decoding schemes such as beam-search or top-K sampling. See Appendix C for more detailed analysis.

4.4 Automatic Evaluation

Metrics: We first used automatic metrics to evaluate our method: 1) *Distinct (Dist)* ([Li et al., 2016](#)) calculates the proportion of unique n-grams (n=1, 2) in the generated responses, which is widely used to measure the response diversity. 2) *Entropy (Ent)* ([Zhang et al., 2018](#)) evaluates how evenly the empirical n-gram (n=1, 2) distribution is. Higher scores mean more diverse of the response. 3) *Low-Frequency Token Ratio (LF)* ([Li et al., 2019](#)) further measures the model diversity by counting the ratio of low-frequency words in the generated responses. We chose words with a frequency less than 100 in each corpus as low-frequency words. Over-confident models tend to omit low-frequency words (i.e., get low **LF** scores) and yield less diversified responses. 4) *BLEU* ([Papineni et al., 2002](#)) measures n-gram (n=2, 3, 4) overlap between the generated responses and references.

Results: As shown in Table 2, our method AdaLabel outperforms all the baselines by large margins on all the datasets. We can further observe that: 1) AdaLabel achieves the best diversity scores (Dist-1,2, Ent-1,2, and LF). This indicates that our method yields better training targets that help to produce more diverse responses; 2). The models that explicitly tackle the over-confidence issue (i.e., AdaLabel and FACE) generally outperform

other baselines in diversity-related metrics. For example, FACE obtains the second-best diversity scores (i.e., Dist, Ent, and LF) on the OpenSubtitles dataset. This verifies our motivation that alleviating the over-confidence issue helps to produce more diverse responses.

Note that our method also outperforms all the baselines using the stochastic decoding scheme. Please refer to Appendix C for more details.

4.5 Manual Evaluation

Metrics: Pairwise manual evaluations are conducted to further validate our method. Specifically, for a given dialogue post, our model’s response is paired with the one from a baseline. Three individual annotators were employed to rank each response pair from three aspects: 1) *Fluency (Flu.)*: which response is more fluent; 2) *Coherency (Coh.)*: which response is more coherent to the context; 3) *Informativeness (Info.)*: which response contains more informative content. We also asked the annotator to choose an overall preferred response (**Pref.**). Ties were allowed.

Results: 200 posts were randomly sampled from each of these two datasets, respectively, and totally 3.6K response pairs were generated. The inter-rater annotation agreement was measured using Fleiss’s kappa κ ([Fleiss, 1971](#)). Particularly, the κ value on DailyDialog, OpenSubtitles dataset was 0.59 and 0.55, respectively, indicating moderate agreement.

As shown in Table 3, AdaLabel outperforms all the baselines on the informativeness measure. This means that our method can respond with more informative content. We can further observe that:

1). All models achieve competitive fluency because it is easy for neural models to produce fluent responses by yielding trivial responses like “I

Comparison	DailyDialog				OpenSubtitles			
	Pref.	Flu.	Coh.	Info.	Pref.	Flu.	Coh.	Info.
AdaLabel vs CE	17.00 [‡]	1.33	12.5 [‡]	28.33 [‡]	6.33	1.17	7.33 [†]	13.67 [‡]
AdaLabel vs LS	2.67	0.17	3.33	24.83 [‡]	5.3	-0.67	3.17	8.50 [‡]
AdaLabel vs FL	4.50	1.67	7.00 [†]	22.0 [‡]	8.00 [†]	1.00	6.00	5.50
AdaLabel vs FACE	6.67 [†]	3.50 [†]	7.17 [†]	8.50 [†]	4.50	0.50	1.83	2.50
AdaLabel vs F ²	7.67 [†]	0.33	6.83 [†]	8.67 [‡]	4.33	-0.50	1.67	9.50 [‡]
AdaLabel vs CP	10.50 [‡]	-0.17	8.00 [†]	23.83 [‡]	8.00 [†]	1.50	6.17	16.83 [‡]
AdaLabel vs UL	7.83 [†]	0.83	6.67 [†]	17.33 [‡]	6.83 [†]	2.00	5.83	15.00 [‡]
AdaLabel vs NL	9.17 [†]	2.67 [†]	9.17 [†]	7.67 [†]	5.17	0.17	2.17	15.5 [‡]
AdaLabel vs D2GPo	0.83	0.00	3.33	15.17 [‡]	3.17	7.33 [‡]	1.00	6.33 [†]

Table 3: Pairwise human evaluation results (%). The absolute gains of AdaLabel (i.e., $Win\ rate - Lose\ rate$) are reported. [†], [‡] indicates significant improvement with p -value < 0.05 and < 0.005 , respectively (sign test).

Model	BLEU-3,4		Dist-1,2		Ent-1,2		LF
1.w/o ε	5.46	3.57	2.52	13.21	4.64	6.89	4.85
2.w/o α	11.35	8.70	3.62	20.56	5.02	7.70	7.30
3.Orig. v	8.15	5.77	3.71	19.53	5.00	7.58	8.25
4.Uniform	5.66	3.61	2.24	14.96	4.84	7.33	4.98
5.Rand	6.27	4.07	2.03	13.47	4.7	7.08	4.56
6.BERT	11.6	9.34	3.67	20.97	5.02	7.71	7.28
AdaLabel	13.38	11.01	3.96	23.53	5.17	8.00	8.49

Table 4: Ablation study results on DailyDialog (%).

don’t know”. However, our model surpasses most baselines in terms of fluency while ensuring high diversity scores. This demonstrates the superiority of our method in producing high quality responses.

2). AdaLabel produces more coherent responses comparing to most baselines. This verifies that our model does not sacrifice the response quality when achieving high diversity scores. In fact, by controlling the model’s confidence, more low-frequency words are encouraged, and thus AdaLabel can produce more relevant and coherent responses. This claim is further verified by observing that our model achieves the best overall preference score among all the baselines.

4.6 Ablation study

Ablation studies were performed to verify the effect of each component in our method. Specifically, two groups of variants were tested:

The first group validates the effectiveness of the calculated target word probability, i.e., ε : **1).** **w/o ε** directly sets a fixed value for ε in Eq. 2. The specific value of ε is searched from 0.1 to 0.7 with a stride of 0.1; **2).** **w/o α** omits the empirical factor α in calculating ε , i.e., the value of ε in Eq. 2 is calculated using Eq. 4 in instead of Eq. 6.

The second group validates the effectiveness of

the non-target word probabilities produced by D_a , i.e., v : **3).** **Orig.** v does not truncate the head of v when inferring from D_a . Note that the truncation for the tail of v is still applied since its effectiveness has already been proved in previous studies (Tang et al., 2020; Tan et al., 2019); **4).** **Uniform** uses an uniform distribution as v in Eq. 2. Note that different from the baseline **LS**, the value of ε is calculated using Eq. 6 in this ablation model, whereas the value of ε in the baseline **LS** is fixed ; **5).** **Rand** use a random distributions as v in Eq. 2; **6).** **BERT** follows the work of Chen et al. (2020) to fine-tune a pre-trained BERT model to produce v . Note that our dialogue model may benefit from the multi-task training of D_a since D_a shares the same encoder with our dialogue model. Optimizing Eq. 8 may help the encoder to capture better features. For fair comparison, we kept the task of optimizing D_a in ablation models 4-6 although it is not used to infer v .

Table 4 shows the results of ablation models on the DailyDialog dataset. As can be seen from the first two rows, our method to adaptively calculate ε helps to improve the performance of our model by a large margin, and the empirical adjustment factor α helps to further improve our performance by facilitating the learning of low-probability words. The performance of ablation models 3-6 in Table 4 proves that v captures reliable distribution and helps our model produce more diverse responses. Moreover, truncating the head distribution of v enables the dialogue model to focus more on the low-frequency words and thus facilitates more informative responses.

It is also interesting to note that our auxiliary decoder D_a surpasses the BERT teacher used by Chen et al. (2020) in helping the dialogue model

	DailyDialog	OpenSubtitles
Auxiliary Decoder D_a	64.03	64.92
Dialog Decoder in AdaLabel	44.16	43.90
Dialog Decoder in CE	38.58	41.57

Table 5: Prediction accuracy of decoders on test sets.

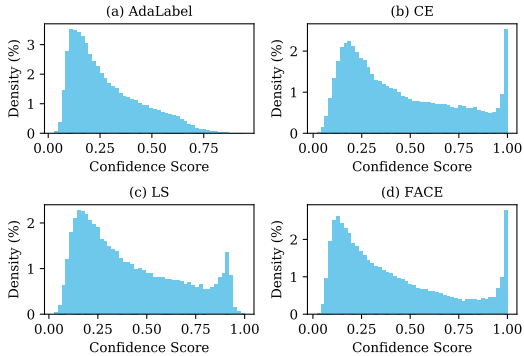


Figure 4: Empirical distribution of confidence scores for high-frequency words on the OpenSubtitles dataset. Words occupying the top 40% of the frequency mass in the training set are regarded as high-frequency words.

to produce more diverse responses. This further proves the effectiveness of D_a considering that BERT contains 6 times parameters than D_a and consumes much more computation resources.

5 Discussion

5.1 Auxiliary Decoder

To further test the performance of D_a , we evaluated the averaged accuracy score of D_a when predicting each target word in the test set (first row in Table 5). Specifically, a target word y_t in the reference response is determined to be correctly predicted if it is top-ranked in the predicted distribution $p(\cdot|y_{<t}, X)$. A better decoder is generally believed to obtain a higher accuracy. Table 5 also reports the uni-directional dialogue decoders’ accuracy in AdaLabel and CE. It can be seen that D_a can make substantially more accurate predictions with the help of modeling bi-directional contexts using only one layer. Moreover, the dialogue model’s decoder in AdaLabel, which is guided by D_a , achieves better accuracies than the CE. This further proves that our light-weight D_a is capable of producing effective v .

5.2 Prediction Confidence

We also visualized the distribution of confidence scores assigned by each dialogue model to high-frequency words. Figure 4 shows the results of

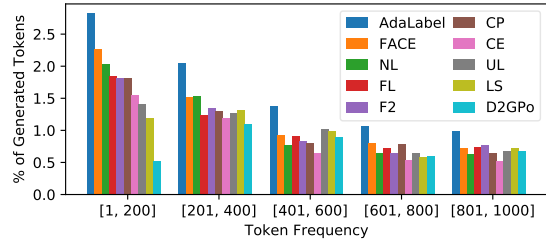


Figure 5: Ratios of low-frequency tokens in the generated responses on the OpenSubtitles dataset. Tokens in each group are determined based on the frequency on the training set.

four best performing models on the OpenSubtitles dataset. The spikes of high confidence score observed in Figure 4b and 4d indicate that CE and FACE assign extremely high confidence scores to a large number of high-frequency words. Although the smoothed labels in LS manage to alleviate these high-confidence-spikes (Figure 4c), a considerable amount of words still receives high confidence scores in LS. Our model outperforms all the baselines to avoid assigning over-confidence scores, thus alleviating the over-confidence issue. A similar trend is also observed on the DailyDialog dataset (see Appendix D for results of all models on both datasets).

5.3 Predicted Rare Word Distribution

Over-confident models produce less diversified responses because they usually under-estimate rare words. To evaluate the effectiveness of AdaLabel, we tested whether AdaLabel encourages more “rare words” in its generations. Specifically, the ratio of generated tokens corresponding to different token frequency bins is calculated, and the results on the OpenSubtitles dataset are shown in Figure 5. It can be seen that AdaLabel produces more rare words in the generated responses than other baselines. Similar results are also observed on the DailyDialog dataset (see Appendix E).

6 Conclusion

We address the low-diversity issue of neural dialogue models by introducing an adaptive label smoothing approach, AdaLabel. In our method, the probability of each target word is estimated based on the current dialogue model’s prediction, and the probabilities for these non-target words are calculated using a novel auxiliary decoder D_a . A target-masked attention scheme is introduced in D_a

to help capture forward and backward contexts. We evaluate our method on two benchmark datasets: DailyDialog and OpenSubtitles. Extensive experiments show that our method effectively alleviates the over-confidence issue and improves the diversity of the generated responses. As future work, we believe this method is extensible to other text generation tasks.

Acknowledgments

This work was partly supported by the NSFC projects (Key project with No. 61936010 and regular project with No. 61876096). This work was also supported by the Guoqiang Institute of Tsinghua University, with Grant No. 2019GQG1 and 2020GQG0005. We thank Jinchao Zhang and Yao Qiu for early discussions and insightful comments of this work.

References

- Hengyi Cai, Hongshen Chen, Yonghao Song, Cheng Zhang, Xiaofang Zhao, and Dawei Yin. 2020. [Data manipulation: Towards effective instance learning for neural dialogue generation via learning to augment and reweight](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 6334–6343, Online. Association for Computational Linguistics.
- Yen-Chun Chen, Zhe Gan, Yu Cheng, Jingzhou Liu, and Jingjing Liu. 2020. [Distilling knowledge learned in BERT for text generation](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 7893–7905, Online. Association for Computational Linguistics.
- Byung-Ju Choi, Jimin Hong, David Park, and Sang Wan Lee. 2020. [F²-softmax: Diversifying neural text generation via frequency factorized softmax](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 9167–9182, Online. Association for Computational Linguistics.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. [BERT: Pre-training of deep bidirectional transformers for language understanding](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4171–4186, Minneapolis, Minnesota. Association for Computational Linguistics.
- Jiacheng Du, Wenjie Li, Yulan He, Ruifeng Xu, Lidong Bing, and Xuan Wang. 2018. [Variational autoregressive decoder for neural response generation](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 3154–3163, Brussels, Belgium. Association for Computational Linguistics.
- Joseph L Fleiss. 1971. Measuring nominal scale agreement among many raters. *Psychological bulletin*, 76(5):378.
- Xiang Gao, Sungjin Lee, Yizhe Zhang, Chris Brockett, Michel Galley, Jianfeng Gao, and Bill Dolan. 2019. [Jointly optimizing diversity and relevance in neural response generation](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 1229–1238, Minneapolis, Minnesota. Association for Computational Linguistics.
- Thamme Gowda and Jonathan May. 2020. [Finding the optimal vocabulary size for neural machine translation](#). In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 3955–3964, Online. Association for Computational Linguistics.
- Chuan Guo, Geoff Pleiss, Yu Sun, and Kilian Q Weinberger. 2017. On calibration of modern neural networks. In *International Conference on Machine Learning*, pages 1321–1330. PMLR.
- Tianxing He and James Glass. 2020. [Negative training for neural dialogue response generation](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 2044–2058, Online. Association for Computational Linguistics.
- Geoffrey E. Hinton, Oriol Vinyals, and Jeffrey Dean. 2015. [Distilling the knowledge in a neural network](#). *CoRR*, abs/1503.02531.
- Ari Holtzman, Jan Buys, Li Du, Maxwell Forbes, and Yejin Choi. 2020. [The curious case of neural text degeneration](#). In *8th International Conference on Learning Representations, ICLR 2020, Addis Ababa, Ethiopia, April 26-30, 2020*. OpenReview.net.
- Minlie Huang, Xiaoyan Zhu, and Jianfeng Gao. 2020. Challenges in building intelligent open-domain dialog systems. *ACM Transactions on Information Systems (TOIS)*, 38(3):1–32.
- Daphne Ippolito, Reno Kriz, João Sedoc, Maria Kustikova, and Chris Callison-Burch. 2019. [Comparison of diverse decoding methods from conditional language models](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 3752–3762, Florence, Italy. Association for Computational Linguistics.
- Shaojie Jiang, Pengjie Ren, Christof Monz, and Maarten de Rijke. 2019. Improving neural response diversity with frequency-aware cross-entropy loss. In *The World Wide Web Conference*, pages 2879–2885.

- Shaojie Jiang and Maarten de Rijke. 2018. [Why are sequence-to-sequence models so dull? understanding the low-diversity problem of chatbots.](#) In *Proceedings of the 2018 EMNLP Workshop SCAI: The 2nd International Workshop on Search-Oriented Conversational AI*, pages 81–86, Brussels, Belgium. Association for Computational Linguistics.
- Jungo Kasai, Nikolaos Pappas, Hao Peng, James Cross, and Noah Smith. 2021. [Deep encoder, shallow decoder: Reevaluating non-autoregressive machine translation.](#) In *International Conference on Learning Representations*.
- Yoon Kim and Alexander M. Rush. 2016. [Sequence-level knowledge distillation.](#) In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 1317–1327, Austin, Texas. Association for Computational Linguistics.
- Diederik P. Kingma and Jimmy Ba. 2015. [Adam: A method for stochastic optimization.](#) In *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*.
- Guillaume Klein, Yoon Kim, Yuntian Deng, Jean Senellart, and Alexander Rush. 2017. [OpenNMT: Open-source toolkit for neural machine translation.](#) In *Proceedings of ACL 2017, System Demonstrations*, pages 67–72, Vancouver, Canada. Association for Computational Linguistics.
- Iliia Kulikov, Alexander Miller, Kyunghyun Cho, and Jason Weston. 2019. [Importance of search and evaluation strategies in neural dialogue modeling.](#) In *Proceedings of the 12th International Conference on Natural Language Generation*, pages 76–87, Tokyo, Japan. Association for Computational Linguistics.
- Aviral Kumar and Sunita Sarawagi. 2019. [Calibration of encoder decoder models for neural machine translation.](#) *CoRR*, abs/1903.00802.
- Jiwei Li, Michel Galley, Chris Brockett, Jianfeng Gao, and Bill Dolan. 2016. [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.
- Margaret Li, Stephen Roller, Iliia Kulikov, Sean Welleck, Y-Lan Boureau, Kyunghyun Cho, and Jason Weston. 2020. [Don't say that! making inconsistent dialogue unlikely with unlikelihood training.](#) In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 4715–4728, Online. Association for Computational Linguistics.
- Yanran Li, Hui Su, Xiaoyu Shen, Wenjie Li, Ziqiang Cao, and Shuzi Niu. 2017. [DailyDialog: A manually labelled multi-turn dialogue dataset.](#) In *Proceedings of the Eighth International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 986–995, Taipei, Taiwan. Asian Federation of Natural Language Processing.
- Zuchao Li, Rui Wang, Kehai Chen, Masso Utiyama, Eiichiro Sumita, Zhuosheng Zhang, and Hai Zhao. 2019. [Data-dependent gaussian prior objective for language generation.](#) In *International Conference on Learning Representations*.
- Tsung-Yi Lin, Priya Goyal, Ross Girshick, Kaiming He, and Piotr Dollár. 2017. [Focal loss for dense object detection.](#) In *Proceedings of the IEEE international conference on computer vision*, pages 2980–2988.
- Luca Massarelli, Fabio Petroni, Aleksandra Piktus, Myle Ott, Tim Rocktäschel, Vassilis Plachouras, Fabrizio Silvestri, and Sebastian Riedel. 2020. [How decoding strategies affect the verifiability of generated text.](#) In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 223–235, Online. Association for Computational Linguistics.
- Jishnu Mukhoti, Viveka Kulharia, Amartya Sanyal, Stuart Golodetz, Philip H. S. Torr, and Puneet K. Dokania. 2020. [Calibrating deep neural networks using focal loss.](#) In *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual*.
- Rafael Müller, Simon Kornblith, and Geoffrey E Hinton. 2019. [When does label smoothing help?](#) In *Advances in Neural Information Processing Systems*, pages 4694–4703.
- Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. 2002. [Bleu: a method for automatic evaluation of machine translation.](#) In *Proceedings of the 40th Annual Meeting of the Association for Computational Linguistics*, pages 311–318, Philadelphia, Pennsylvania, USA. Association for Computational Linguistics.
- Gabriel Pereyra, George Tucker, Jan Chorowski, Lukasz Kaiser, and Geoffrey E. Hinton. 2017. [Regularizing neural networks by penalizing confident output distributions.](#) In *5th International Conference on Learning Representations, ICLR 2017, Toulon, France, April 24-26, 2017, Workshop Track Proceedings*. OpenReview.net.
- Christian Szegedy, Vincent Vanhoucke, Sergey Ioffe, Jon Shlens, and Zbigniew Wojna. 2016. [Rethinking the inception architecture for computer vision.](#) In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2818–2826.
- Xu Tan, Yi Ren, Di He, Tao Qin, Zhou Zhao, and Tie-Yan Liu. 2019. [Multilingual neural machine translation with knowledge distillation.](#) In *7th International Conference on Learning Representations, ICLR 2019, New Orleans, LA, USA, May 6-9, 2019*. OpenReview.net.

- Jiayi Tang, Rakesh Shivanna, Zhe Zhao, Dong Lin, Anima Singh, Ed H. Chi, and Sagar Jain. 2020. [Understanding and improving knowledge distillation](#). *CoRR*, abs/2002.03532.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin. 2017. [Attention is all you need](#). In *Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems 2017, December 4-9, 2017, Long Beach, CA, USA*, pages 5998–6008.
- Yida Wang, Pei Ke, Yinhe Zheng, Kaili Huang, Yong Jiang, Xiaoyan Zhu, and Minlie Huang. 2020. A large-scale chinese short-text conversation dataset. In *CCF International Conference on Natural Language Processing and Chinese Computing*, pages 91–103. Springer.
- Sean Welleck, Ilya Kulikov, Stephen Roller, Emily Dinan, Kyunghyun Cho, and Jason Weston. 2020. [Neural text generation with unlikelihood training](#). In *8th International Conference on Learning Representations, ICLR 2020, Addis Ababa, Ethiopia, April 26-30, 2020*. OpenReview.net.
- Chenglin Yang, Lingxi Xie, Siyuan Qiao, and Alan L. Yuille. 2018. [Knowledge distillation in generations: More tolerant teachers educate better students](#). *CoRR*, abs/1805.05551.
- Rongsheng Zhang, Yinhe Zheng, Jianzhi Shao, Xiaoxi Mao, Yadong Xi, and Minlie Huang. 2020. Dialogue distillation: Open-domain dialogue augmentation using unpaired data. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 3449–3460.
- Yizhe Zhang, Michel Galley, Jianfeng Gao, Zhe Gan, Xiujuan 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 1810–1820.
- Tiancheng Zhao, Ran Zhao, and Maxine Eskenazi. 2017. [Learning discourse-level diversity for neural dialog models using conditional variational autoencoders](#). In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 654–664, Vancouver, Canada. Association for Computational Linguistics.
- Yinhe Zheng, Zikai Chen, Rongsheng Zhang, Shilei Huang, Xiaoxi Mao, and Minlie Huang. 2020a. [Stylized dialogue response generation using stylized unpaired texts](#). In *AAAI*.
- Yinhe Zheng, Rongsheng Zhang, Minlie Huang, and Xiaoxi Mao. 2020b. A pre-training based personalized dialogue generation model with persona-sparse data. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, pages 9693–9700.
- Ganbin Zhou, Ping Luo, Yijun Xiao, Fen Lin, Bo Chen, and Qing He. 2018. Elastic responding machine for dialog generation with dynamically mechanism selecting. In *Thirty-Second AAAI Conference on Artificial Intelligence*.

A Implementation Details

This appendix describes the implementation details of our model. All our experiments are implemented with python 3.7.4, PyTorch 1.7.1, and the OpenNMT package (Klein et al., 2017). Training is performed on one TITAN Xp GPU. Our model’s backbone is the transformer-based sequence to sequence model, the encoder and decoder each contains 6 transformer layers with 8 attention heads, and the hidden size is set to 512. The dimension of the feed-forward layer is also 512. The WordPiece tokenizer provided by BERT-base-uncased is used (the vocabulary contains 30522 tokens). The total number of parameters in our model is about 90M. The Adam optimizer is employed to train our model from random initializations with $\beta_1 = 0.9$, $\beta_2 = 0.999$, $\epsilon = 1e - 9$ and a learning rate of $1e-4$. The batch size is set to 64 with 2 gradient accumulation so that $2 * 64$ samples are used for each parameter update. The model is evaluated every 1000 steps on the validation set. We use early-stopping with patience 10, 30 for DailyDialog and OpenSubtitles, respectively. Specifically, the model stops training when the evaluation perplexity and accuracy are not increased for “patience” steps. The model training takes 4 hours and 3 days on DailyDialog and OpenSubtitles, respectively.

The auxiliary distribution produced by the auxiliary decoder is smoothed with the temperature scaling approach. The temperature used in this process is searched in [1, 1.5, 2]. The temperature value of 1.5 and 1.0 is used for DailyDialog, and OpenSubtitles, respectively. The hyper-parameter value of η is set to 0.2 for all datasets. The fixed value of epsilon in our ablation model **w/o** ϵ is searched in [0.1, 0.2, 0.3, 0.4, 0.5, 0.6], and we find the value of 0.1 works best.

B Baseline Implementation Details

This appendix contains more implementation details of our baselines. All the baselines utilize the same backbone architecture and basic hyperparameter settings as our model (see Appendix A). The hyper-parameters specialized for each baseline is determined with the grid search based on the Dist

Model	DailyDialog								OpenSubtitles							
	Dist-1, 2		Ent-1, 2		LF	BLEU-2,3,4			Dist-1, 2		Ent-1, 2		LF	BLEU-2,3,4		
CE	1.79	8.21	4.19	5.90	2.57	4.06	2.49	1.58	2.48	9.21	4.07	5.74	0.76	7.03	4.26	2.82
LS	1.71	8.01	4.16	5.89	2.17	4.13	2.55	1.65	2.89	12.79	4.27	6.24	0.47	8.24	5.57	4.20
FL	2.40	11.37	4.39	6.35	4.46	6.01	3.95	2.75	3.10	12.37	4.25	6.13	0.82	7.13	4.56	3.25
FACE	1.80	9.47	4.54	6.40	3.48	5.65	3.43	2.17	3.12	12.62	4.47	6.40	1.02	5.97	3.63	2.43
F ²	1.61	7.22	4.04	5.70	2.11	4.32	2.55	1.52	2.89	10.63	4.03	5.72	0.89	6.92	4.27	2.91
CP	2.30	10.39	4.28	6.16	3.25	5.31	3.39	2.30	3.14	11.87	4.17	5.97	0.85	7.28	4.60	3.21
UL	2.42	11.0	4.40	6.42	4.55	7.94	5.26	3.69	2.77	10.43	3.98	5.62	0.62	6.89	4.36	3.03
NL	1.61	7.53	4.19	6.05	4.02	7.09	4.41	2.91	2.65	10.14	4.21	6.05	0.75	7.16	4.32	2.85
D2GPo	1.57	7.83	4.14	5.91	2.26	4.47	2.71	1.71	2.06	10.43	4.15	6.00	0.12	7.32	4.69	3.33
AdaLabel	4.25	21.47	4.95	7.51	7.68	14.71	11.63	9.80	4.91	21.53	4.71	7.08	1.35	8.68	6.08	4.68
AdaLabel (Greedy)	3.96	23.53	5.17	8.00	8.49	17.42	13.38	11.01	4.78	22.88	4.96	7.66	1.47	9.80	6.48	4.75
Human	6.59	37.74	5.67	8.91	13.7	N/A	N/A	N/A	8.62	43.16	5.89	9.36	4.75	N/A	N/A	N/A

Table 6: Automatic evaluation results (%) using the beam search decoding scheme (beam size is 5). The best results among all these beam-search-decoded models are in bold.

measures on the validation set: For **Label smoothing (LS)**, we searched the smoothing parameter in [0.05, 0.1, 0.2, 0.3, 0.4, 0.5], and found 0.1 works best on all the datasets; For **Confidence penalty (CP)**, we searched the weight of penalty in [0.0005, 0.001, 0.01, 0.05, 0.1] and found 0.05 works best on all the datasets while ensuring the loss to be positive; For **Focal loss (FL)**, we searched the hyperparameter γ in [0.1, 0.5, 1, 2, 3], and found 2 works best on all the datasets. For **Unlikelihood loss (UL)**, we searched the weight of penalty in [1, 10, 100, 1000], and select 1000 on all the datasets. For **FACE**, we experiment with the Output token frequency & PRE-weight version, which is reported to be the best version of FACE. For **Negative loss (NL)**, **F²-softmax (F²)** and **Data-dependent Gaussian Prior objective (D2GPo)**, the selection of hyper-parameters follows the author’s suggestion.

C Automatic Evaluation Results with Other Decoding Schemes

This appendix reports our model’s automatic evaluation results and all the baselines when different decoding schemes are used. Specifically, Table 6 shows the results for the beam search decoding scheme (beam size of 5), and Table 7 shows the results when the top-K decoding scheme ($k = 10$) is used. Note that for the F²-softmax, we use the decoupled top-k sampling as the authors suggested.

As can be seen from Table 6 and 7, our method outperforms all the baselines on the diversity-related scores (i.e., Dist, Ent, and LF) by a large margin. This indicates that our method can produce

more diverse responses even with the stochastic based decoding scheme.

We also include the results of AdaLabel when the greedy decoding scheme is used in Table 6 and Table 7 (the second line from the bottom). It is interesting to see that the greedily decoded responses from AdaLabel are more diverse than some baselines that are decoded using the sampling scheme (see Table 7). Moreover, our model AdaLabel with the greedy decoding scheme achieves the best BLEU among all the baselines on both datasets.

D Prediction Confidence

This appendix reports the prediction confidence scores assigned by each model to high-frequency words. Specifically, words occupying the top 40% of the frequency mass in the training set of each dataset are regarded as high-frequency words.

Figure 6 shows the results of our model and all the baselines on the DailyDialog dataset. Figure 7 shows the results of our model and all the baselines on the OpenSubtitles dataset. It can be seen that most of our baselines assign extremely high confidence scores (nearly 1.0) to these high-frequency words, and thus resulting in a spike of high confidence scores in the plotted distribution. Our model outperforms all the baselines in avoiding assigning extremely high confidence scores to these high-frequency words.

E Predicted Rare Word Distribution on DailyDialog

This appendix shows the distribution of rare words in the generated responses on the DailyDialog

Model	DailyDialog								OpenSubtitles							
	Dist-1, 2		Ent-1, 2		LF	BLEU-2,3,4			Dist-1, 2		Ent-1, 2		LF	BLEU-2,3,4		
CE	2.22	19.05	5.07	7.87	4.09	6.78	3.29	1.61	3.78	20.58	5.07	7.97	1.23	5.94	2.84	1.46
LS	1.95	17.74	5.02	7.82	3.69	7.08	3.50	1.77	3.46	21.27	5.10	8.12	0.78	6.15	3.16	1.85
FL	2.71	20.98	5.19	8.17	6.44	8.09	4.13	2.24	3.82	22.14	5.15	8.25	1.27	5.34	2.54	1.34
FACE	2.29	21.14	5.36	8.3	5.73	7.07	3.47	1.82	4.25	23.95	5.30	8.37	1.51	5.34	2.54	1.33
F ²	2.16	19.33	5.04	7.85	3.97	6.31	3.12	1.58	4.10	22.53	5.13	8.11	1.32	5.27	2.51	1.31
CP	3.16	22.38	5.11	7.96	6.01	8.11	4.38	2.50	4.06	22.62	5.13	8.14	1.33	6.00	2.94	1.52
UL	2.92	20.81	5.12	7.99	6.44	9.36	5.13	3.00	3.74	20.97	5.01	7.94	1.00	6.01	2.99	1.64
NL	2.39	18.35	4.99	7.79	5.72	8.71	4.64	2.63	3.57	20.36	5.05	7.97	1.06	5.84	2.86	1.46
D2GPo	1.75	17.09	5.00	7.81	3.40	7.45	3.73	1.97	2.74	19.21	5.00	7.97	0.36	6.32	3.15	1.72
AdaLabel	4.11	32.65	5.58	8.93	10.99	8.87	4.84	2.90	4.78	29.58	5.43	8.78	1.53	5.12	2.32	1.19
AdaLabel (Greedy)	3.96	23.53	5.17	8.00	8.49	17.42	13.38	11.01	4.78	22.88	4.96	7.66	1.47	9.80	6.48	4.75
Human	6.59	37.74	5.67	8.91	13.7	N/A	N/A	N/A	8.62	43.16	5.89	9.36	4.75	N/A	N/A	N/A

Table 7: Automatic evaluation results (%) using the top-k sampling decoding scheme ($k = 10$). The best results among all these top-k-decoded models are in bold.

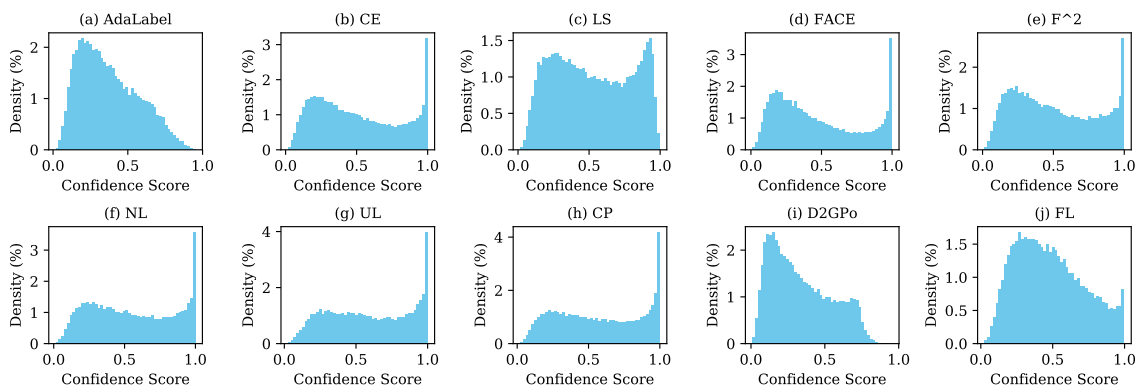


Figure 6: Confidence score distributions for high-frequency words on the DailyDialog dataset. Words occupying the top 40% of the frequency mass in the training set of DailyDialog are regarded as high-frequency words.

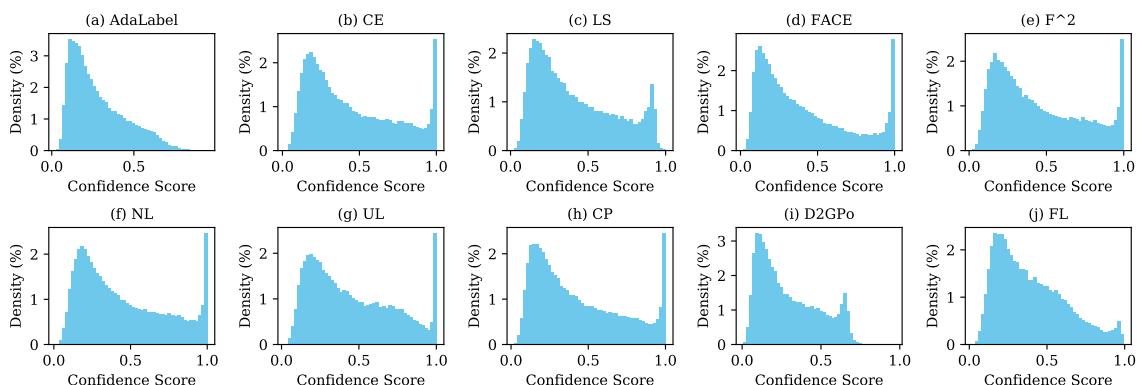


Figure 7: Confidence score distributions for high-frequency words on the OpenSubtitles dataset. Words occupying the top 40% of the frequency mass in the training set of OpenSubtitles are regarded as high-frequency words.

dataset (see Figure 8). It can be seen that more “rare words” are predicted by our method on the DailyDialog dataset. This observation is in line with the results on the OpenSubtitles dataset as reported in Section 5.3.

F Use BERT Model to Obtain v

This appendix provides more experiment results comparing to the CMLM model (Chen et al., 2020): **1). CMLM** exactly follows the setting of Chen et al. (2020), i.e., the teacher distribution produced by

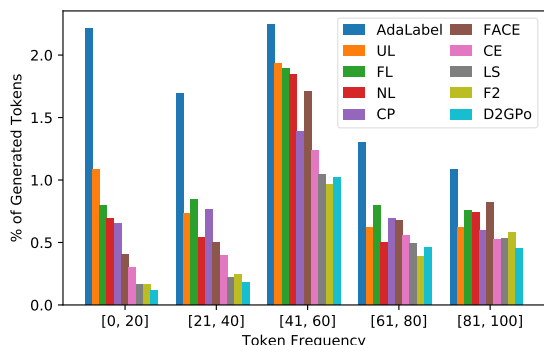


Figure 8: Ratios of low-frequency tokens in the generated responses on the DailyDialog dataset. Tokens in each group are determined based on the frequency on the training set.

Model	BLEU-3,4	Dist-1,2	Ent-1,2	LF			
1. CMLM	6.18	4.09	2.20	11.83	4.59	6.79	4.62
2. CMLM+ ϵ	9.36	7.31	3.78	21.05	4.96	7.61	6.88
3. CMLM+ ϵ + D_a	11.6	9.34	3.67	20.97	5.02	7.71	7.28
AdaLabel	13.38	11.01	3.96	23.53	5.17	8.00	8.49

Table 8: Ablation study results based on BERT on DailyDialog (%).

the BERT model is merged with the one-hot distribution using a fixed ϵ . **2).** CMLM+ ϵ adaptively adjust the value of ϵ using Eq. 6 in our paper. **3).** CMLM+ ϵ + D_a add an additional training task to optimize the auxiliary decoder D_a on the basis of CMLM+ ϵ . It is expected that optimizing D_a help our dialogue encoder to capture better representations. The trained D_a is not used in the training and inference phase of our dialogue model. Note that the last model CMLM+ ϵ + D_a is the same with our ablation model **6. BERT** as reported in our paper.

As can be seen Table 8, our approach to adaptively change ϵ helps to produce better dialogue responses, and the training of D_a helps our dialogue encoder to learn better representations.

G Case study

We sampled some generated cases on the DailyDialog and OpenSubtitles dataset. The results of our model and some competitive baselines are shown in Table 9 and Table 10. It can be seen that the responses generated by our method are coherent to the context and contain richer contents. Moreover, our model also produces more rare words that make our response more diverse.

Dialogue	
Context1	Tell me a little bit about yourself , please .
UL	I am sorry , sir . I have been in a foreign currency company for the company .
D2GPo	I ' m afraid i ca n ' t . I ' m a little nervous about it .
FACE	Yes , i am in the mood for a few days ago .
AdaLabel	I am studying in the international trade company . I am studying english in my company .
Context1	Oh , they look nice . may i try that black one on ?
Context2	Of course . the fitting room is right over there .
UL	I think you ' re right . i think you ' ll have a try .
D2GPo	I think it ' s a good idea .
FACE	How much is it ?
AdaLabel	I bought this one . it ' s the latest fashion .
Context1	Well , we don ' t have any pre - existing illnesses or conditions , so we should be fine .
UL	And most people are n ' t so quiet and interviewing .
D2GPo	And the company company are a good job in the company .
FACE	And your company insurance , or workshops are booked in a private company ?
AdaLabel	I wish our company or country provided us with healthcare .

Table 9: Cases on DailyDialog.

Dialogue	
Context	You ' re going to love this baby .
LS	I ' m going to get you .
D2GPo	I ' m going to go to the bathroom .
FACE	You ' re going to love this baby .
AdaLabel	You ' re going to be a good father .
Context	Mr. Matlack can ' t offend .
LS	This is our official business .
D2GPo	This is the most important thing .
FACE	It ' s a pleasure to meet you , mr . President .
AdaLabel	Timothy Matlack is the official scribe of the continental congress .
Context	Brave men discover they are really cowards .
LS	I ' m a samurai .
D2GPo	I ' ve never seen such a beast in my life .
FACE	They are not afraid of you .
AdaLabel	Father , you ' re a coward !
Context	I found this whole operation strange .
LS	I ' m not going to let you go .
D2GPo	I ' m not sure i ' m going to do it .
FACE	I ' m sorry , but i ' m not sure i can handle it .
AdaLabel	It ' s an incredibly delicate system .
Context	If they make it , they ' re clear into a safe zone where they can get medical supplies and food .
LS	We ' il get them to the safe .
D2GPo	We ' il have to get back to the hotel .
FACE	They ' re gon na get us out of here .
AdaLabel	So we can use it as a safe field .

Table 10: Cases on OpenSubtitles.