

Distilling the Knowledge of Large-scale Generative Models into Retrieval Models for Efficient Open-domain Conversation

Beomsu Kim[†], Seokjun Seo[†], Seungju Han[†], Enkhbayar Erdenee[†], Buru Chang^{*}
Hyperconnect

{beomsu.kim, seokjun.seo, seungju.han, enkhbayar.erdenee, buru.chang}@hpcnt.com

Abstract

Despite the remarkable performance of large-scale generative models in open-domain conversation, they are known to be less practical for building real-time conversation systems due to high latency. On the other hand, retrieval models could return responses with much lower latency but show inferior performance to the large-scale generative models since the conversation quality is bounded by the pre-defined response set. To take advantage of both approaches, we propose a new training method called G2R (Generative-to-Retrieval distillation) that preserves the efficiency of a retrieval model while leveraging the conversational ability of a large-scale generative model by infusing the knowledge of the generative model into the retrieval model. G2R consists of two distinct techniques of distillation: the data-level G2R augments the dialogue dataset with additional responses generated by the large-scale generative model, and the model-level G2R transfers the response quality score assessed by the generative model to the score of the retrieval model by the knowledge distillation loss. Through extensive experiments including human evaluation, we demonstrate that our retrieval-based conversation system trained with G2R shows a substantially improved performance compared to the baseline retrieval model while showing significantly lower inference latency than the large-scale generative models.

1 Introduction

Recently, generative models have shown great success in open-domain conversation along with the development of large-scale language models, yielding fluent and informative responses (Roller et al., 2021; Adiwardana et al., 2020; Brown et al., 2020). However, generative models suffer from the challenges of latency and computational resources for

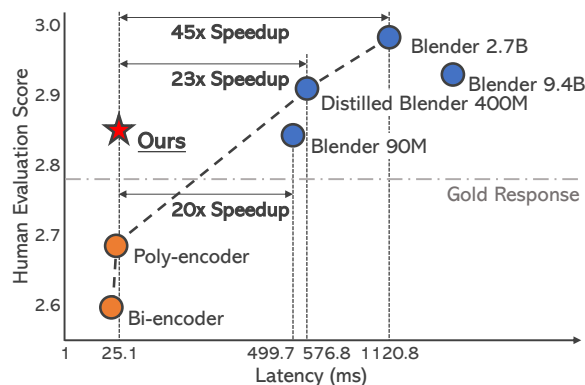


Figure 1: Latency vs. Human evaluation score plot for open-domain conversation models. Blue circle represents generative models, orange circle represents retrieval models, and red star represents our model. Our model achieves a "sweet-spot" among various models, showing a much better human evaluation score than retrieval models and demonstrating much lower latency than generative models.

building real-time conversation systems due to auto-regressive decoding for response generation and a large GPU memory footprint.

Meanwhile, retrieval models such as Bi-encoder and Poly-encoder (Humeau et al., 2019) is able to build efficient open-domain conversation systems by pre-defining the response set and searching the most relevant response to the given context from the response set. In addition, a Bi-encoder dramatically reduces the latency when adopting efficient Maximum Inner Product Search (MIPS) libraries, such as FAISS (Johnson et al., 2019) and ScaNN (Guo et al., 2020). Despite the outstanding efficiency, retrieval models have shown some lack of conversational ability compared to generative models. Retrieval models are known to return an erroneous response when the pre-defined response set does not contain the proper response to the given context, while generative models deal with these cases more flexibly (Weston et al., 2018). Exemplar-based generative models (Weston et al., 2018; Wu et al., 2019; Gupta et al., 2021) try to

[†]Equal contribution

^{*}Corresponding author

mitigate this problem by combining the advantages of the two approaches, whereas the inherent inefficiency of the generative models remains since exemplar-based generative models employ a generative model for response generation.

To make an efficient yet fluent open-domain conversation system, which is mandatory for real-world applications, we propose a novel training method for retrieval models called **Generative-to-Retrieval distillation** (G2R). G2R enables retrieval models to leverage the knowledge of large-scale generative models in both data-level and model-level. First, **data-level G2R** augments the original dialogue dataset with the responses produced by a large-scale generative model using contexts in the original dialogue dataset. Then, the produced responses are also added to the pre-defined response set. The augmented dialogue dataset and response set are utilized for training a retrieval model at the training phase and for returning responses at the inference phase, respectively. Although data-level G2R enables retrieval model to utilize high-quality responses generated by the large-scale generative model, it does not transfer the fine-grained knowledge from the generative model about the quality of individual responses. **Model-level G2R** resolves this limitation by transferring the response quality scores assessed by the large-scale teacher generative model into the scores of the student retrieval model. This method induces the retrieval model to select a better response in terms of the response quality.

We empirically demonstrate that a retrieval-based conversation system, which consists of the G2R-applied retrieval model and a MIPS library, shows a substantial conversational ability while showing fast inference speed, as shown in Figure 1. For instance, our retrieval-based conversation system shows about a 20x speedup compared to the Blender model (90M parameters) while exhibiting a comparable human evaluation result on conversational ability.

2 Method

2.1 Preliminaries

Retrieval models for Open-domain Conversation. Let $D = \{(c_i, r_i) \mid 1 \leq i \leq n\}$ denote the dialogue dataset that contains n context-response pairs, where c_i and r_i are a context and its corresponding gold response of the i -th example, respectively. At the training phase, retrieval models are

trained to maximize the score of the gold response r_i for the given context c_i compared to the scores of negative responses. At the inference phase, retrieval models return the response with the highest score for the given context c from the pre-defined response set $R = \{r_i \mid 1 \leq i \leq n\}$ constructed from the dialogue dataset D .

Knowledge Distillation. Knowledge Distillation (Hinton et al., 2015) transfers the knowledge of the teacher model into the student model by adding a loss that matches the logits of the student model z_s with the logits of the teacher model z_t . For classification task with l classes, the knowledge distillation loss is defined by the cross-entropy between the softened output probability of the student model and the teacher model:

$$\begin{aligned} \mathcal{L}_{KD} &= - \sum_{x \in X} \sum_{i=1}^l p_t(y_i|x) \log p_s(y_i|x) \\ &= - \sum_{x \in X} \sum_{i=1}^l \left[\frac{\exp(z_t(x, y_i)/T)}{\sum_j \exp(z_t(x, y_j)/T)} \times \right. \\ &\quad \left. \log \frac{\exp(z_s(x, y_i)/T)}{\sum_j \exp(z_s(x, y_j)/T)} \right], \quad (1) \end{aligned}$$

where $p(y|x)$ and $z(x, y)$ are the softened probability and logit value of the models for the input x and class y , respectively, and T is a temperature parameter for smoothing the logit values.

2.2 Retrieval-based Conversation System

Our goal is to create an efficient open-domain conversation system based on the retrieval model. However, naively utilizing the retrieval model can lead to the low efficiency when the size of the response set R is large since the retrieval model has to calculate scores for all response candidates. To this end, we adopt the Bi-encoder (Humeau et al., 2019) model with an efficient MIPS library to select proper responses efficiently without calculating a score for all response candidates. Bi-encoder encodes a context c and response r into the fixed-length embedding respectively with Transformer architecture (Vaswani et al., 2017), and defines the relevance score between c and r as the dot-product of two embeddings. Therefore, an efficient MIPS library, FAISS (Johnson et al., 2019) for our case, can be utilized for speeding up the search process.

2.3 Data-level G2R

It is well-known that utilizing an additional high-quality dialogue dataset is helpful for improving

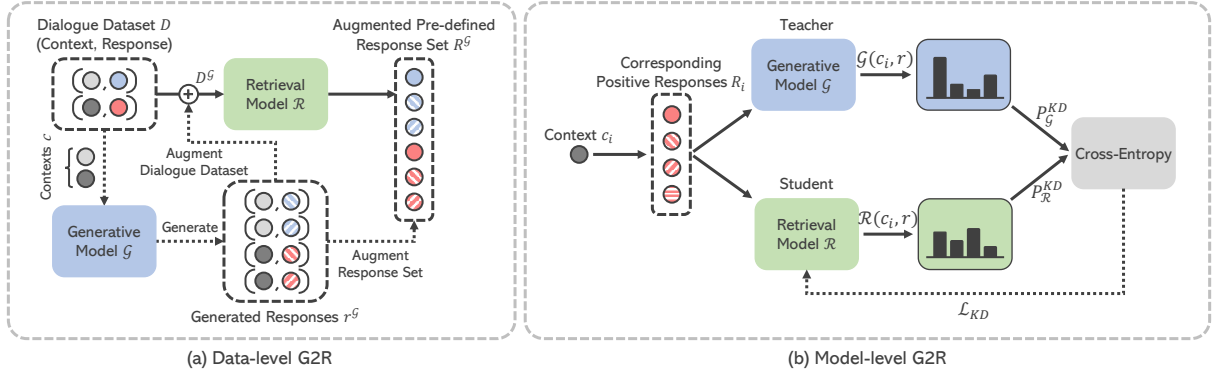


Figure 2: Illustration of our proposed training method G2R.

the performance of the retrieval model, as shown in Zhang et al. (2020). Moreover, enriching the pre-defined response set R with more diverse responses can help the retrieval model to respond appropriately to a variety of input contexts since it widens the opportunity to select an appropriate response. However, it is highly labor-intensive and costly to acquire such high-quality dialogue datasets or responses through human-in-the-loop annotation such as in Zhang et al. (2018) or Smith et al. (2020).

Meanwhile, previous studies (Adiwardana et al., 2020; Roller et al., 2021; Brown et al., 2020) show that well-tuned large-scale generative models are able to achieve near-human conversational ability. From these observations, we are motivated to leverage the generation result of large-scale generative models to extend the response set as well as the dialogue dataset for training a retrieval model, as shown in Figure 2(a).

For each context c_i in the dialogue dataset D , a large-scale generative model \mathcal{G} generates m responses, $\{r_{i,j}^{\mathcal{G}} \mid 1 \leq j \leq m\}$. Considering the generated responses as a gold response of the given context c_i , they are added to the dialogue dataset D and the pre-defined response set R as follows: $D^{\mathcal{G}} = D \cup \{(c_i, r_{i,j}^{\mathcal{G}}) \mid 1 \leq i \leq n, 1 \leq j \leq m\}$ and $R^{\mathcal{G}} = R \cup \{r_{i,j}^{\mathcal{G}} \mid 1 \leq i \leq n, 1 \leq j \leq m\}$. $D^{\mathcal{G}}$ and $R^{\mathcal{G}}$ denote the augmented dialogue dataset and response set, respectively.

After the augmentation, a retrieval model \mathcal{R} is trained by minimizing the cross-entropy loss \mathcal{L}_{CE} which maximizes the probability of selecting the ground-truth response r among the set of randomly sampled negative responses R^- :

$$\mathcal{L}_{CE} = - \sum_{(c,r) \in D^{\mathcal{G}}} \log \frac{\exp(\mathcal{R}(c, r))}{\sum_{r^- \in \{r\} \cup R^-} \exp(\mathcal{R}(c, r^-))}, \quad (2)$$

where $\mathcal{R}(c, r)$ is the score computed by the retrieval model \mathcal{R} for the given context c and response r . Note that R^- is created differently for every iteration by randomly sampling responses from $R^{\mathcal{G}}$ without replacement.

We employ the largest open-domain conversation model available, Blender 9.4B (Roller et al., 2021), as the large-scale generative model \mathcal{G} . We apply top-k sampling (Fan et al., 2018) for the diversity of responses since beam search tends to generate similar responses within the same context (Adiwardana et al., 2020). In addition, we sample responses multiple times with different minimum length constraints to diversify the specificity and length of generated responses.

2.4 Model-level G2R

While data-level G2R provides additional high-quality dialogue data and diverse responses, it does not transfer the fine-grained knowledge about the quality of the individual responses from the large-scale generative model \mathcal{G} . Model-level G2R is designed to address this problem by transferring the individual response-level quality score, assessed by the large-scale teacher generative model \mathcal{G} , into the student retrieval model \mathcal{R} . We first define the quality score of the response from the perspective of the teacher generative model \mathcal{G} , denoted as $\mathcal{G}(c, r)$. Then, the student retrieval model is trained to match the score $\mathcal{R}(c, r)$ of the student retrieval model with the score $\mathcal{G}(c, r)$ of the teacher generative model, similar to the conventional knowledge distillation technique (Hinton et al., 2015). Overall process of knowledge distillation is depicted in Figure 2(b).

We define the generator score $\mathcal{G}(c, r)$ as the log-likelihood normalized by the length of response:

$$\mathcal{G}(c, r) = (\log P_{\mathcal{G}}(r|c))/|r|, \quad (3)$$

where $P_{\mathcal{G}}(r|c)$ is the probability of the response r for the given context c of the generative model \mathcal{G} and $|r|$ is the number of tokens in the response r . Log-likelihood is normalized with the length of response to mitigate the problem of preferring shorter responses (Murray and Chiang, 2018).

We can derive the distillation loss \mathcal{L}_{KD} by regarding the generator quality score $\mathcal{G}(c, r)$ and retriever score $\mathcal{R}(c, r)$ as the logits of teacher and student model, respectively. Eq. 1 then turns into:

$$\begin{aligned} P_{\mathcal{G}}^{KD}(c_i, r) &= \frac{\exp(\mathcal{G}(c_i, r)/T)}{\sum_{r' \in R_i \cup R^-} \exp(\mathcal{G}(c_i, r')/T)}, \\ P_{\mathcal{R}}^{KD}(c_i, r) &= \frac{\exp(\mathcal{R}(c_i, r)/T)}{\sum_{r' \in R_i \cup R^-} \exp(\mathcal{R}(c_i, r')/T)}, \\ \mathcal{L}_{KD} &= - \sum_{i=1}^n \sum_{r \in R_i \cup R^-} P_{\mathcal{G}}^{KD}(c_i, r) \log P_{\mathcal{R}}^{KD}(c_i, r), \end{aligned} \quad (4)$$

where $R_i = \{r_i, r_{i,1}^{\mathcal{G}}, \dots, r_{i,m}^{\mathcal{G}}\}$ is a set of positive responses correspond to the context c_i in $D^{\mathcal{G}}$. Since calculating the generator quality score $\mathcal{G}(c_i, r^-)$ for negative responses requires heavy extra computation, we simplify the calculation by approximating $P_{\mathcal{G}}^{KD}(c_i, r^-) \approx 0$, $\exp(\mathcal{G}(c_i, r^-)/T) \approx 0$ for randomly sampled negative responses $r^- \in R^-$.

Our final loss for the model-level G2R is a sum of original cross-entropy loss in Equation 2 and the knowledge distillation loss where hyperparameter α controls the weights of each term:

$$\mathcal{L} = \alpha \mathcal{L}_{CE} + (1 - \alpha) \mathcal{L}_{KD}. \quad (5)$$

3 Experiments

3.1 Dataset

We conduct experiments on the open-domain conversation datasets which consist of Blended Skill Talk (Smith et al., 2020), ConvAI2 (Zhang et al., 2018), Empathetic Dialogues (Rashkin et al., 2019) and Wizard of Wikipedia (Dinan et al., 2018). Following Roller et al. (2021), all four datasets are used together for the following experiments, and we refer to the merged dataset as BST+. We follow the method of splitting train, validation, and test set from Smith et al. (2020).

3.2 Metrics

Human Evaluation. We conduct a human evaluation to assess the quality of model responses. Human evaluation is carried out on 200 examples

randomly sampled from the BST+ test dataset. Human judges are asked to evaluate the quality of the generated response with two criteria on a 0-2 scale: (i) *Appropriateness (Appr.)* for evaluating whether the generated response is fluent, logical, and appropriate to its given context, and (ii) *Informativeness (Info.)* for evaluating whether the generated response has meaningful information relevant to its given context. Each example is rated by at least three unique human judges, and all the human evaluation is performed via Amazon Mechanical Turk.

Automated Metrics. We also report various kinds of automated metrics. *MaUdE* (Sinha et al., 2020) is an unreferenced dialogue response evaluation metric calculated by the model that is trained to score positive responses as 1 while scoring syntactically and semantically negative responses as 0, using the ConvAI2 dataset. Since *MaUdE* shows a high correlation with human judgments on fluency and interestingness of responses, we use *MaUdE* as a proxy metric for evaluating the overall quality of responses produced by each model. For measuring the lexical diversity of generated responses we utilize *Dist-2* and *Dist-3* (Li et al., 2016), where *Dist-n* is a ratio of distinct n-grams to the total number of n-grams in all the responses generated by each model. *Length*, the average number of tokens in generated responses, is reported for reference. Last but not least, we measure and report the *Latency* for generating a response for a single input context to verify the efficiency of the model. Although we report the latency measured on the GPU-enabled environment, the latency measured by using only the CPU is reported in the supplementary material.

3.3 Models and Baselines

Blender. Blender, the state-of-the-art model in open-domain conversation task, is adopted with different number of parameters: *Blender 90M*, *Blender 2.7B*, and *Blender 9.4B*. For response generation, we follow the decoding hyperparameters suggested in the original work (Roller et al., 2021).

Distilled Blender. A small Blender model distilled from a larger generative model is employed to compare our result with a generative model that also utilizes the knowledge distillation technique. We use 400M parameters Blender model distilled from *Blender 2.7B* with TinyBERT style distillation (Jiao et al., 2020), denoted as *Distilled Blender*.

Bi-encoder & Poly-encoder. *Bi-encoder* and *Poly-*

Models	Human Evaluation			Automated Metrics				Latency	Latency
	Sum	Appr.	Info.	MaUdE	Dist-2	Dist-3	Length	(ms)	(Speedup)
Blender 90M	2.843	1.429	1.414	0.8582	0.4799	0.6887	18.31	499.7	1.00x
Blender 2.7B	2.983	1.510	1.473	0.8826	0.5288	0.7261	19.05	1120.8	0.45x
Blender 9.4B	2.930	1.472	1.458	0.8763	0.5246	0.7285	18.87	1438.6	0.35x
Distilled Blender	2.910	1.474	1.436	0.8715	0.4821	0.6815	19.19	576.8	0.87x
RetNRef	2.771	1.404	1.368	0.8555	0.7773	0.9541	12.34	382.4	1.31x
Bi-encoder	2.597	1.288	1.309	0.8944	0.8191	0.9712	14.85	18.6	26.87x
Poly-encoder	2.686	1.340	1.346	0.8645	0.8269	0.9692	15.30	24.8	20.15x
Bi-encoder (w/ FAISS)	2.596	1.259	1.337	0.9046	0.8316	0.9735	15.22	25.7	19.44x
G2R-D (w/o FAISS)	2.779	1.380	1.399	0.8518	0.7242	0.9302	20.06	39.7	12.59x
G2R-D	2.759	1.398	1.361	0.8443	0.7456	0.9395	19.93	25.3	19.75x
G2R-DM	2.856	1.447	1.410	0.8695	0.7266	0.9393	17.48	25.1	19.91x
Human Response	2.788	1.418	1.369	0.9146	0.8271	0.9742	14.22	-	-

Table 1: Human evaluation results and automated metrics of the baseline models and our G2R models. *Latency (Speedup)* column denotes the relative speedup of each model compared to the latency of *Blender 90M*.

encoder with 256M parameters (Humeau et al., 2019), pre-trained with the Pushshift Reddit comment dataset (Baumgartner et al., 2020) and fine-tuned on the BST+ dataset, are the baselines for retrieval models. The *Bi-encoder* model integrated with MIPS library, as described in Section 2.2, is denoted as *Bi-encoder (w/ FAISS)*.

RetNRef. *RetNRef* (Weston et al., 2018) is an exemplar-based generative model which incorporates the response of retrieval models into the input of the generative model. Contrary to G2R, *RetNRef* exploits the retrieval model to make the generative model better, while G2R exploits the knowledge of the generative model to make the retrieval model better. We use the dialogue retrieval model described in Roller et al. (2021) trained with the α -blending technique.

Human Response. *Human response* refers to the ground-truth label annotated in the BST+ dataset.

G2R. Our system is built upon the retrieval-based conversation system described in Section 2.2, where the Bi-encoder \mathcal{R} is trained with our proposed G2R using *Blender 9.4B* as the teacher generative model \mathcal{G} . *G2R-DM* denotes our model trained with the data-level G2R and the model-level G2R. For a comprehensive analysis, two variants are adopted: *G2R-D* is trained with the data-level G2R only, and *G2R-D (w/o FAISS)* further removes the use of the MIPS library, FAISS, from *G2R-D*.

3.4 Implementation Details

We provide the details on our implementation and the hyperparameter values in the supplementary material. For reproducibility, we release the augmented dialogue dataset and the implementation of

G2R models.¹

4 Experimental Results

4.1 Result Analysis

We present the human evaluation result and the automated metrics in Table 1. Overall, our system trained with G2R achieves a "sweet-spot" between conversational ability and efficiency. Our system maintains the low latency of *Bi-encoder (w/ FAISS)* while boosting up the human evaluation results significantly, achieving comparable or better human evaluation scores than the *Blender 90M* and human responses, respectively.

Taking a closer look, the Blender generative models and the distilled variant show high human evaluation metric while showing relatively large latency along with the lack of diversity, as shown in the Dist-2 and Dist-3 scores. Retrieval baselines (*Bi-encoder* and *Poly-encoder*) show an opposite trend, exhibiting much lower latency and relatively higher response diversity but showing relatively lower conversational ability in terms of human evaluation metric. Unlike human evaluation results, the MaUdE scores of the *Bi-encoder* and the *Poly-Encoder* are unexpectedly high. However, we suspect this is because the MaUdE metric is trained on the ConvAI2 dataset, which is a subset of the BST+ dataset, and with a similar training objective of these retrieval models as described in Section 3.

G2R-based models achieve far better human evaluation results compared to their original model, *Bi-encoder (w/ FAISS)*. Applying data-level G2R only (*G2R-D*) significantly boosts the performance,

¹<https://github.com/hyperconnect/g2r>

Statistics	R	R^G	Ratio
# of Responses	279,090	3,070,074	11.0x
Average length	14.85	18.78	1.26x
# of Unique Tokens	56,862	210,538	3.70x
# of Unique bi-grams	655,948	2,710,155	4.13x
# of Unique tri-grams	1,738,189	10,654,181	6.13x

Table 2: Comparison of the statistics of the original response set R and the response set R^G augmented by data-level G2R. We also report the ratio of the statistics of the augmented dataset to those of the original dataset.

making the model perform comparable to gold human response in terms of human evaluation. Using data-level G2R enlarges the number of responses in the pre-defined response set R^G more than ten times, therefore using Bi-encoder without FAISS (*G2R-D (w/o FAISS)*) leads to increased latency. Although using FAISS induces a latency overhead for the case where the size of the response set is small (*Bi-encoder (w/ FAISS)*), using FAISS in a larger response set as in *G2R-D* enables us to maintain the low latency, while having a slight degradation of response qualities compared to the version without FAISS.

Further application of model-level G2R additionally boosts the performance of the retrieval model. *G2R-DM* shows a higher human evaluation score and MaUdE score than *G2R-D*, and exhibits a comparable human evaluation score to the *Blender 90M* model while running much faster. While *G2R-DM* shows a somewhat deficient human evaluation score compared to the bigger Blender generative models, it shows substantially lower latency (23.0x speedup over *Distilled Blender*, 44.7x speedup over *Blender 2.7B*). In addition, *G2R-DM* exhibits a much higher response diversity compared to the Blender generative models. The *RetNRef* model shows worse performance and delivers much higher latency compared to our *G2R-DM* model.

4.2 Statistics of the Responses augmented by the Data-level G2R

Table 2 shows the basic statistics of the original response set R and the response set R^G created by data-level G2R. After applying the data-level G2R, R^G has roughly 11 times more candidates compared to the original response set R . To verify if responses in the new response set R^G show more diversity, we count the number of unique tokens and bi-gram/tri-grams appearing in each response set. The augmented response set has much

more unique tokens and bi-gram/tri-grams than the original response set, implying that it covers more diverse topics, entities and shows more diversity in terms of phrases and expressions.

4.3 Ablation Studies

Breakdown analysis of Data-level G2R. We conduct an ablation study to analyze in detail how the performance of the model changes depending on how we use responses generated in the data-level G2R method. In data-level G2R, generated responses are utilized in two ways: for augmenting the training dialogue dataset D^G of the retrieval model \mathcal{R} , and for building the augmented response set R^G . We separate these two utilization methods and evaluate models that use only each method.

Table 3 shows the evaluation results of these ablation models. Along with the human evaluation metrics and automated metrics, we also report *Hits@1/K* and *Hits@5/K* (Roller et al., 2021) of trained Bi-encoder model on the BST+ test set, which are widely adopted to evaluate the performance of retrieval models. As shown in the table, only utilizing one of the methods does not show better performance compared to the model that utilizes both methods. Utilizing the generated responses for building R^G improves the appropriateness score of the model, which supports the hypothesis we have raised in Section 2 that using a diverse response set is helpful for the model to respond more appropriately. The use of augmented dialogue D^G for training \mathcal{R} is helpful for increasing a human evaluation score, for both appropriateness and informativeness metrics, meaning that the retrieval model learns to select relevant and rich responses that the generative model created. In addition, training with augmented dialogue D^G considerably improves the Hits metric of the retrieval model. Nonetheless, using both methods shows the best human evaluation performance among all ablation models, indicating that using new examples for both training a retrieval model and building a response set is crucial for inducing a good performance.

Different Dialogue Augmentation Strategy. Here, we implement a simple baseline inspired by Zhu et al. (2020) and Zhang et al. (2020), which augments training dialogue by utilizing top- m responses of a retrieval model that has already been trained. In this experiment, we use the *Bi-encoder* model for this augmentation process, and the augmented dialogue dataset generated by this method

Train \mathcal{R} with	Response Set	Human Evaluation			Automated Metrics				
		Sum	Appr.	Info.	Dist-2	Dist-3	Length	Hits@1/K	Hits@5/K
D	R	2.596	1.259	1.337	0.8336	0.9777	15.66	0.7537	0.9363
	$R^{\mathcal{G}}$	2.620	1.300	1.320	0.7660	0.9498	17.14		
$D^{\mathcal{G}}$	R	2.739	1.377	1.361	0.8144	0.9687	16.20	0.8052	0.9570
	$R^{\mathcal{G}}$	2.770	1.403	1.368	0.7456	0.9395	19.93		
$D^{\mathcal{R}}$	R	2.591	1.296	1.295	0.8253	0.9669	14.54	0.7594	0.9323

Table 3: Human evaluation and automated metric results of the ablation models for data-level G2R. Note that $D^{\mathcal{R}}$ is inspired from [Zhu et al. \(2020\)](#), and is not the G2R method.

$\mathcal{G}(c, r)$	Human Evaluation			Automated Metrics		
	Sum	Appr.	Info.	MaUdE	Dist-2	Dist-3
LL	2.856	1.447	1.410	0.8695	0.7266	0.9393
MI	2.806	1.427	1.380	0.8737	0.7536	0.9468

Table 4: Human evaluation results and automated metrics for model-level G2R models that use different score for defining generator quality score $\mathcal{G}(c, r)$.

is denoted as $D^{\mathcal{R}}$. Comparison of data-level G2R with this baseline will enable us to verify that our method with a large generative model produces better quality training dataset than simply using a retrieval model.

The result of this ablation study is reported in Table 3. As shown in the table, using $D^{\mathcal{R}}$ as the training dataset does not lead to a significant performance gain for all metrics, contrary to the case of using $D^{\mathcal{G}}$ which improves both human evaluation score and Hits metric. This result strongly indicates that utilizing a large-scale generative model for dialogue augmentation as in data-level G2R is a much more effective augmentation strategy than using retrieval models.

Utilizing a Different Generator Quality Score for Model-level G2R. Although we employ the log-likelihood score (LL score) for defining the generator quality score $\mathcal{G}(c, r)$ in model-level G2R, there are other methods that can be utilized as well. One example is a Mutual Information score (MI score) ([Li et al., 2016](#)), which is a point-wise mutual information between the given context c and response r . Details about calculating the MI score of response are described in the supplementary material. MI score is known to assign lower values to generic responses while escalating the score of responses that are more specific to the given context, so we expect that using the MI score will produce more specific and diverse responses compared to the LL score while having a slightly higher risk of returning responses with inappropriate details with

respect to input context. Therefore, we evaluate the variant of model-level G2R that uses MI score as $\mathcal{G}(c, r)$ and compare the performance with the model that uses LL score.

The results are provided in Table 4. Using MI score for the model-level G2R exhibits a slightly lower human evaluation score than using LL score, especially for the appropriateness score, implying that using the MI score may be less capable of producing an appropriate and accurate answer. However, in terms of automated metrics, the MI score shows a higher MaUdE score. In addition, using MI score shows higher response diversity compared to LL score as expected, indicating that MI score could be employed for more diverse responses of the conversation system.

4.4 Case Study

Table 5 provides an example of responses returned by the baseline models and our G2R models. In this example, *Bi-encoder (w/ FAISS)* returns the irrelevant response to the given context. Blender models’ responses are logically appropriate, however, they just simply change the topic (*Blender 90M, Distilled Blender*) or relatively lack of specific details (*Blender 2.7B, Blender 9.4B*). *G2R-D* tries to respond with detail, but the response contains a somewhat irrelevant phrase about groceries. In contrast, *G2R-DM* respond appropriately along with specific details talking about a particular book title. We provide additional response examples in the supplementary material.

5 Related Work

5.1 Open-domain Conversation

The task of open-domain conversation has been studied based on retrieval models, generation models, or using both. While retrieval models ([Wang et al., 2013](#); [Ji et al., 2014](#); [Wang et al., 2015](#); [Yan et al., 2016](#); [Wu et al., 2017](#); [Zhou et al., 2018](#); [Tao](#)

Input Context	
A: Amazon is a great place to order books B: Oh yeah that is a fact.	
Model Name	Response
Bi-encoder (w/ FAISS)	A: Comics is at the top of the list in interest
Blender 90M	A: Do you have any other hobbies that you like to do on the weekends? I like to read
Distilled Blender	A: What else do you like to do in your spare time? I like to go to the beach
Blender 2.7B	A: What kind of books do you buy on amazon? I like to read science fiction and fantasy
Blender 9.4B	A: What kind of books do you like to read? I like romance novels and mystery novels
G2R-D	A: Do you ever order groceries? I love amazon's selection of books and videos
G2R-DM	A: What is your favourite book? Mine is "the cat in the hat" by dr seuss

Table 5: Example responses in the BST+ test set example. Full dialogue context is shown in the supplementary material.

et al., 2019; Humeau et al., 2019) search a response relevant to a given context from a pre-defined response set, generative models (Shang et al., 2015; Vinyals and Le, 2015; Li et al., 2020; Holtzman et al., 2019; Welleck et al., 2019; Roller et al., 2021) produce a response based on the given context with auto-regressive decoding. It is known that the retrieval and generative models have advantages in the efficiency of inference and quality of generated responses, respectively. To take both advantages, several exemplar-based generative models (Guu et al., 2018; Wu et al., 2019; Weston et al., 2018; Cai et al., 2019; Gupta et al., 2021) have recently been proposed by combining the retrieval and generative models. The main difference between our proposed training method and the exemplar-based generative models is that exemplar-based generative models provide the knowledge of retrieval models to generative models, while our proposed training method transfers the knowledge of generative models to retrieval models to focus on the efficiency of open-domain conversation systems.

5.2 Knowledge Transfer from Large Models

Transferring the knowledge from larger-scale teacher neural networks into smaller-scale student neural networks has been implemented to improve the performance of the student model, including data augmentation and knowledge distillation. In the data augmentation perspective, several works (Schick and Schütze, 2021; Chang et al., 2021; Kumar et al., 2020; Yang et al., 2020) utilize the generation result of pre-trained language models as a labeled example for text classification tasks. Lin et al. (2020b) utilize the inference result of the retrieval model and the generative model as a semi-negative dataset for training a student retrieval model. Meanwhile, Knowledge distillation

(Hinton et al., 2015) transfers the knowledge of the teacher model into the student model by matching the student logits with softened teacher logits. Knowledge distillation especially designed for specific tasks or model architectures exists, such as sequence generation task (Kim and Rush, 2016; Lin et al., 2020a), retrieval models (Lu et al., 2020; Vakili Tahami et al., 2020) and for transformer architectures (Jiao et al., 2020; Wang et al., 2020; Sun et al., 2020).

The most related work to our paper is Dialogue Distillation (Zhang et al., 2020), which also proposes a data-level and model-level distillation for open-domain conversation models. Our research differs from this work in three ways. First, Dialogue Distillation requires additional unpaired text corpus, which could be hard to be obtained in certain circumstances. We instead focus on utilizing the knowledge of large-scale generative models for augmenting additional data. In addition, Dialogue Distillation does not enrich the pre-defined response set, which is crucial for improving the performance of the retrieval models, as shown in our experiments. Last but not least, while Dialogue Distillation only considers the distillation within the homogeneous architecture, Generative-to-Generative or Retrieval-to-Retrieval, we focus on the model-level distillation between heterogeneous architectures, especially Generative-to-Retrieval, to take advantages of each architecture.

6 Conclusion

We present G2R, a novel training scheme of retrieval model for open-domain conversation by distilling the knowledge of large-scale generative models in both data-level and model-level. G2R enables retrieval models to build a highly efficient conversation system that exhibits a substantial level of con-

versational ability. We believe that our work will serve as a stepping stone for creating an efficient and real-time open-domain conversation system.

Ethical Considerations

We train our models with the BST+ dataset, and the models we used for the pre-training (Pre-trained Bi-encoder weights from Humeau et al. (2019)) and generating the augmented dataset (Blender 9.4B) are trained with the Pushshift Comment Dataset (Baumgartner et al., 2020) and the BST+ dataset. Both the BST+ dataset and the Pushshift dataset are publicly available. Texts included in these datasets may include potentially abusive contents and underlying biases, and these toxicities and biases could have been unintentionally encoded in our models. Therefore, methods for reducing the toxicity of the open-domain dialogue system (Xu et al., 2020; Dinan et al., 2019) or methods for mitigating the bias of the dialogue model (Liu et al., 2020; Dinan et al., 2020) are recommended to be jointly used with our method when deploying our model in production.

Like any other open-domain conversational system, our system might provide false or misleading information. Furthermore, our system has the potential to return a response that contains private information. Since our model is a retrieval-based model and the pre-defined response set is fixed, an effort for filtering out the responses that potentially contain false information, private information, profanity, and inappropriate content should be preceded.

We acknowledge that it is possible to have biases in human evaluation through Amazon Mechanical Turk. To reduce potential biases, we set a maximum number of annotations per worker. We did not ask the user’s identity; therefore, their personal information, including their gender, race, ethnicity, etc., is not revealed.

References

- Daniel Adiwardana, Minh-Thang Luong, David R So, Jamie Hall, Noah Fiedel, Romal Thoppilan, Zi Yang, Apoorv Kulshreshtha, Gaurav Nemade, Yifeng Lu, et al. 2020. Towards a human-like open-domain chatbot. *arXiv preprint arXiv:2001.09977*.
- Jason Baumgartner, Savvas Zannettou, Brian Keegan, Megan Squire, and Jeremy Blackburn. 2020. The pushshift reddit dataset. In *Proceedings of the International AAAI Conference on Web and Social Media*, volume 14, pages 830–839.
- Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh, Daniel Ziegler, Jeffrey Wu, Clemens Winter, Chris Hesse, Mark Chen, Eric Sigler, Mateusz Litwin, Scott Gray, Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec Radford, Ilya Sutskever, and Dario Amodei. 2020. *Language models are few-shot learners*. In *Advances in Neural Information Processing Systems*, volume 33, pages 1877–1901. Curran Associates, Inc.
- Deng Cai, Yan Wang, Wei Bi, Zhaopeng Tu, Xiaojiang Liu, and Shuming Shi. 2019. Retrieval-guided dialogue response generation via a matching-to-generation framework. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 1866–1875.
- Ernie Chang, Vera Demberg, and Alex Marin. 2021. Jointly improving language understanding and generation with quality-weighted weak supervision of automatic labeling. In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume*, pages 818–829.
- Emily Dinan, Angela Fan, Adina Williams, Jack Urbanek, Douwe Kiela, and Jason Weston. 2020. Queens are powerful too: Mitigating gender bias in dialogue generation. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 8173–8188.
- Emily Dinan, Samuel Humeau, Bharath Chintagunta, and Jason Weston. 2019. Build it break it fix it for dialogue safety: Robustness from adversarial human attack. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 4537–4546.
- Emily Dinan, Stephen Roller, Kurt Shuster, Angela Fan, Michael Auli, and Jason Weston. 2018. Wizard of wikipedia: Knowledge-powered conversational agents. In *International Conference on Learning Representations*.
- Angela Fan, Mike Lewis, and Yann Dauphin. 2018. Hierarchical neural story generation. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 889–898.
- Ruiqi Guo, Philip Sun, Erik Lindgren, Quan Geng, David Simcha, Felix Chern, and Sanjiv Kumar. 2020. *Accelerating large-scale inference with anisotropic vector quantization*. In *International Conference on Machine Learning*.

- Prakhar Gupta, Jeffrey P Bigham, Yulia Tsvetkov, and Amy Pavel. 2021. Controlling dialogue generation with semantic exemplars. In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 3018–3029.
- Kelvin Guu, Tatsunori B Hashimoto, Yonatan Oren, and Percy Liang. 2018. Generating sentences by editing prototypes. *Transactions of the Association for Computational Linguistics*, 6:437–450.
- Geoffrey Hinton, Oriol Vinyals, and Jeff Dean. 2015. Distilling the knowledge in a neural network. *arXiv preprint arXiv:1503.02531*.
- Ari Holtzman, Jan Buys, Li Du, Maxwell Forbes, and Yejin Choi. 2019. The curious case of neural text degeneration. In *International Conference on Learning Representations*.
- Samuel Humeau, Kurt Shuster, Marie-Anne Lachaux, and Jason Weston. 2019. Poly-encoders: Architectures and pre-training strategies for fast and accurate multi-sentence scoring. In *International Conference on Learning Representations*.
- Zongcheng Ji, Zhengdong Lu, and Hang Li. 2014. An information retrieval approach to short text conversation. *arXiv preprint arXiv:1408.6988*.
- Xiaoqi Jiao, Yichun Yin, Lifeng Shang, Xin Jiang, Xiao Chen, Linlin Li, Fang Wang, and Qun Liu. 2020. Tinybert: Distilling bert for natural language understanding. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: Findings*, pages 4163–4174.
- Jeff Johnson, Matthijs Douze, and Hervé Jégou. 2019. Billion-scale similarity search with gpus. *IEEE Transactions on Big Data*.
- 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.
- Diederik P Kingma and Jimmy Ba. 2014. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*.
- Varun Kumar, Ashutosh Choudhary, and Eunah Cho. 2020. Data augmentation using pre-trained transformer models. In *Proceedings of the 2nd Workshop on Life-long Learning for Spoken Language Systems*, pages 18–26.
- Jiwei Li, Michel Galley, Chris Brockett, Jianfeng Gao, and William B 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.
- Margaret Li, Stephen Roller, Ilya 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.
- Raymond Li, Samira Kahou, Hannes Schulz, Vincent Michalski, Laurent Charlin, and Chris Pal. 2018. Towards deep conversational recommendations. *arXiv preprint arXiv:1812.07617*.
- Alexander Lin, Jeremy Wohlwend, Howard Chen, and Tao Lei. 2020a. Autoregressive knowledge distillation through imitation learning. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 6121–6133.
- Zibo Lin, Deng Cai, Yan Wang, Xiaojiang Liu, Haitao Zheng, and Shuming Shi. 2020b. The world is not binary: Learning to rank with grayscale data for dialogue response selection. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 9220–9229.
- Haochen Liu, Wentao Wang, Yiqi Wang, Hui Liu, Zitao Liu, and Jiliang Tang. 2020. Mitigating gender bias for neural dialogue generation with adversarial learning. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 893–903.
- Edward Loper and Steven Bird. 2002. Nltk: The natural language toolkit. *arXiv preprint cs/0205028*.
- Wenhao Lu, Jian Jiao, and Ruofei Zhang. 2020. Twinbert: Distilling knowledge to twin-structured bert models for efficient retrieval. *arXiv preprint arXiv:2002.06275*.
- Yu A Malkov and Dmitry A Yashunin. 2018. Efficient and robust approximate nearest neighbor search using hierarchical navigable small world graphs. *IEEE transactions on pattern analysis and machine intelligence*, 42(4):824–836.
- Alexander H Miller, Will Feng, Adam Fisch, Jiasen Lu, Dhruv Batra, Antoine Bordes, Devi Parikh, and Jason Weston. 2017. Parlai: A dialog research software platform. *arXiv preprint arXiv:1705.06476*.
- Kenton Murray and David Chiang. 2018. Correcting length bias in neural machine translation. In *Proceedings of the Third Conference on Machine Translation: Research Papers*, pages 212–223.
- Hannah Rashkin, Eric Michael Smith, Margaret Li, and Y-Lan Boureau. 2019. Towards empathetic open-domain conversation models: A new benchmark and dataset. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 5370–5381.

- Stephen Roller, Emily Dinan, Naman Goyal, Da Ju, Mary Williamson, Yinhan Liu, Jing Xu, Myle Ott, Eric Michael Smith, Y-Lan Boureau, et al. 2021. Recipes for building an open-domain chatbot. In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume*, pages 300–325.
- Timo Schick and Hinrich Schütze. 2021. Generating datasets with pretrained language models. *arXiv preprint arXiv:2104.07540*.
- Lifeng Shang, Zhengdong Lu, and Hang Li. 2015. Neural responding machine for short-text conversation. In *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 1577–1586.
- Koustuv Sinha, Prasanna Parthasarathi, Jasmine Wang, Ryan Lowe, William L Hamilton, and Joelle Pineau. 2020. Learning an unreferenced metric for online dialogue evaluation. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 2430–2441.
- Eric Michael Smith, Mary Williamson, Kurt Shuster, Jason Weston, and Y-Lan Boureau. 2020. Can you put it all together: Evaluating conversational agents’ ability to blend skills. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 2021–2030.
- Zhiqing Sun, Hongkun Yu, Xiaodan Song, Renjie Liu, Yiming Yang, and Denny Zhou. 2020. Mobilebert: a compact task-agnostic bert for resource-limited devices. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 2158–2170.
- Chongyang Tao, Wei Wu, Can Xu, Wenpeng Hu, Dongyan Zhao, and Rui Yan. 2019. One time of interaction may not be enough: Go deep with an interaction-over-interaction network for response selection in dialogues. In *Proceedings of the 57th annual meeting of the association for computational linguistics*, pages 1–11.
- Amir Vakili Tahami, Kamyar Ghajar, and Azadeh Shakery. 2020. Distilling knowledge for fast retrieval-based chat-bots. In *Proceedings of the 43rd International ACM SIGIR Conference on Research and Development in Information Retrieval*, pages 2081–2084.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. *Advances in Neural Information Processing Systems*, 30:5998–6008.
- Oriol Vinyals and Quoc Le. 2015. A neural conversational model. *arXiv preprint arXiv:1506.05869*.
- Hao Wang, Zhengdong Lu, Hang Li, and Enhong Chen. 2013. A dataset for research on short-text conversations. In *Proceedings of the 2013 Conference on Empirical Methods in Natural Language Processing*, pages 935–945.
- Mingxuan Wang, Zhengdong Lu, Hang Li, and Qun Liu. 2015. Syntax-based deep matching of short texts. In *Proceedings of the 24th International Conference on Artificial Intelligence*, pages 1354–1361.
- Wenhui Wang, Furu Wei, Li Dong, Hangbo Bao, Nan Yang, and Ming Zhou. 2020. Minilm: Deep self-attention distillation for task-agnostic compression of pre-trained transformers. *arXiv preprint arXiv:2002.10957*.
- Sean Welleck, Ilia Kulikov, Stephen Roller, Emily Dinan, Kyunghyun Cho, and Jason Weston. 2019. Neural text generation with unlikelihood training. In *International Conference on Learning Representations*.
- Jason Weston, Emily Dinan, and Alexander Miller. 2018. Retrieve and refine: Improved sequence generation models for dialogue. In *Proceedings of the 2018 EMNLP Workshop SCAI: The 2nd International Workshop on Search-Oriented Conversational AI*, pages 87–92.
- Yu Wu, Furu Wei, Shaohan Huang, Yunli Wang, Zhoujun Li, and Ming Zhou. 2019. Response generation by context-aware prototype editing. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, pages 7281–7288.
- Yu Wu, Wei Wu, Chen Xing, Ming Zhou, and Zhoujun Li. 2017. Sequential matching network: A new architecture for multi-turn response selection in retrieval-based chatbots. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 496–505.
- Jing Xu, Da Ju, Margaret Li, Y-Lan Boureau, Jason Weston, and Emily Dinan. 2020. Recipes for safety in open-domain chatbots. *arXiv preprint arXiv:2010.07079*.
- Rui Yan, Yiping Song, and Hua Wu. 2016. Learning to respond with deep neural networks for retrieval-based human-computer conversation system. In *Proceedings of the 39th International ACM SIGIR conference on Research and Development in Information Retrieval*, pages 55–64.
- Yiben Yang, Chaitanya Malaviya, Jared Fernandez, Swabha Swayamdipta, Ronan Le Bras, Ji-Ping Wang, Chandra Bhagavatula, Yejin Choi, and Doug Downey. 2020. G-daug: Generative data augmentation for commonsense reasoning. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: Findings*, pages 1008–1025.

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.

Saizheng Zhang, Emily Dinan, Jack Urbanek, Arthur Szlam, Douwe Kiela, and Jason Weston. 2018. Personalizing dialogue agents: I have a dog, do you have pets too? In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 2204–2213.

Xiangyang Zhou, Lu Li, Daxiang Dong, Yi Liu, Ying Chen, Wayne Xin Zhao, Dianhai Yu, and Hua Wu. 2018. Multi-turn response selection for chatbots with deep attention matching network. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1118–1127.

Qingqing Zhu, Xiwei Wang, Chen Chen, and Junfei Liu. 2020. Data augmentation for retrieval-and generation-based dialog systems. In *2020 IEEE 6th International Conference on Computer and Communications (ICCC)*, pages 1716–1720. IEEE.

A Implementation Details

A.1 Baseline models

Blender Models. For Blender models (*Blender 90M*, *Blender 2.7B*, *Blender 9.4B*, *Distilled Blender*), we use the pre-trained weights released from ParlAI (Miller et al., 2017). For a generation, we follow the decoding hyperparameters suggested from the original work (Roller et al., 2021) - using beam search with beam size 10, minimum beam length 20, and tri-gram beam blocking on context and response blocks.

Retrieval Models. We train the *Bi-encoder* and the *Poly-encoder* baseline model on BST+ dataset with pre-trained weights released in ParlAI (Miller et al., 2017), which is originally disclosed by Humeau et al. (2019). Both models have a network parameter size of 256M. We train the model with BST+ dataset, with the batch size 512 and the configuration of using other responses in batch as random negatives, initial learning rate of 1e-5, *ReduceOnPlateau* learning rate schedule with decay rate 0.5 and patience 1. The validation Hits@1/K metric is employed as a proxy metric. Also, we utilize Adamax optimizer (Kingma and Ba, 2014) with gradient clip value 0.1 for our experiments. Note that most of the hyperparameters follow the default implementation of Humeau et al. (2019) implemented in the ParlAI library. These learning

hyperparameters were also used for training other retrieval models in this paper, unless stated.

RetNRef. We train the RetNRef model with a 256M Bi-encoder model architecture as a retriever and 90M Blender generative model architecture as a generator. We follow the α -blending training scheme of (Roller et al., 2021), using blending parameter $\alpha = 0.5$. The model was trained with a batch size of 32 and an initial learning rate of 7e-6, with *ReduceOnPlateau* learning rate scheduler with validation PPL as a proxy metric (with decay rate 0.5, patience 1). For inference, we use the same decoding hyperparameters as in Blender generative models except for the minimum beam length constraint parameter. We used 0 for this value since using a larger value induced a severe repeating problem in the generated response and hurt the performance of the model.

A.2 FAISS

FAISS (Johnson et al., 2019) is employed as an efficient MIPS library for our retrieval-based conversation pipeline. Hierarchical Navigable Small World approximation (Malkov and Yashunin, 2018) is used for building a FAISS index, which was empirically found to be fast and accurate. We use *HNSW32_Flat* index with *efSearch* parameter 256 while using FAISS throughout our implementation.

A.3 Data-level G2R

We use the Blender 9B model (Roller et al., 2021) as our large-scale generative model \mathcal{G} . Throughout our experiments, we use the BST+ training dataset as the original dialogue dataset D , without using the meta-information such as the persona information from ConvAI2 (Zhang et al., 2018) and the Wikipedia topic information from WoW (Dinan et al., 2018). We use top-k sampling with $k = 20$ and tri-gram beam blocking on context and response blocks. We sample 5 samples each from two configurations that use the beam min length hyperparameter of 10 and 20, respectively, sampling a total of 10 samples from a single context c_i . We mainly used ParlAI (Miller et al., 2017) for our experiments. For training Data-level G2R based retrieval model, we compose a mini-batch by randomly selecting 48 unique contexts and randomly selecting 10 responses connected to each context, resulting in a total of 480 (context, response) pairs in a single batch. 512 random negatives are uniformly sampled from response repository $R^{\mathcal{G}}$ and used as a shared random negative

among the examples in the batch. We use the Bi-encoder model trained for baseline retrieval model as initial weights and use the same learning configuration as in the baseline retrieval model except for the initial learning rate value of $5e-5$. We tested the initial learning rate value $lr \in \{1e-5, 5e-5\}$ and selected $5e-5$ since this value has shown faster convergence and higher validation Hits@1/K metric. We trained the model until the convergence of validation Hits@1/K metric and chose the model with the best Hits@1/K metric along the training process. Training takes about 16 to 24 hours in a single NVIDIA DGX Station A100 workstation.

A.4 Model-level G2R

For model-level G2R, we use hyperparameter of $\alpha = 0.9$, and $T = 1$. We did not perform hyperparameter search on T , and tried $\alpha \in \{0.5, 0.9\}$ and selected 0.9 for α since 0.9 has shown higher validation Hits@1/K metric. We use the same training configuration as we train the model in data-level G2R.

B Metrics Details

B.1 Human evaluation

For accurate human evaluation, we only received an annotation from turkers that satisfies the following requirements: (1) HITs approval rate greater than 95%, (2) Location is one of Australia, Canada, New Zealand, United Kingdom, and the United States, (3) Lifetime number of HITs approved greater than 1000, following Li et al. (2018). The instruction for the human evaluation is provided below:

Given the dialogue context, you need to rate the quality of the given response in terms of **appropriateness** and **informativeness**.

Appropriateness is a metric for evaluating whether **the given response is fluent, logical, and appropriate to its given context**. Please rate appropriateness with the range of 0 to 2, where 0 represents bad, and 2 represents excellent. Assign a lower score to the response if the response seems off (illogical, out of context, confusing).

Informativeness is a metric for evaluating whether **the given response has meaningful information relevant to its**

given context. Please rate informativeness with the range of 0 to 2, where 0 represents bad, and 2 represents excellent. Please assign a higher score if the response is rich and specific to the context and a lower score if the response is bland and generic.

B.2 Measuring Latency

We use NVIDIA DGX Station A100 for measuring the latency of the model, with Pytorch 1.7.1, Cuda 11.0, CuDNN 8.0. We only utilize a single GPU (NVIDIA A100 GPU, 40GB Memory) for measuring the latency. Latency is measured as the average inference time of 200 response generations after having 3 warm-up steps.

B.3 Details for Calculating Metrics

For calculating the MaUdE (Sinha et al., 2020) metric, we used the code provided by the authors². For calculating the Dist-2, Dist-3 metrics, and Length, we tokenized the generated response with the `casual_tokenize` method of the `nltk` library (Loper and Bird, 2002) and calculated the metric over 200 generated responses.

C CPU Latency

Model Name	Latency (ms)
Bi-encoder	145.5
G2R-DM-LL (w/o FAISS)	419.4
G2R-DM-LL	163.9
Blender 90M	1908.0
Distilled Blender	9295.8

Table 6: Latency of the models measured by only using CPU.

In Table 6, we report the latency of various models measured by using only CPU. While retrieval models, especially *Bi-encoder* and *G2R-DM*, show an acceptable latency under 200ms, generation models such as *Blender 90M* and *Distilled Blender* exhibit inordinately high latency over 1 second. In particular, *Distilled Blender* shows the latency of 9.3 seconds. The immensely high latency of generative models makes it extremely difficult to employ these models to build real-time conversation agents

²https://github.com/facebookresearch/online_dialog_eval

Model Name	Valid Hits@1/K	Valid Hits@5/K	Test Hits@1/K	Test Hits@5/K
Bi-encoder	0.7469	0.9280	0.7537	0.9363
Bi-encoder \mathcal{R} in G2R-D	0.8011	0.9559	0.8052	0.9570
Bi-encoder \mathcal{R} in G2R-DM	0.8011	0.9558	0.8043	0.9601

Table 7: Hits@N/K metrics of G2R models measured on the validation and the test split of BST+ dataset.

Input Context c_i			
A: Hi I used to be a butcher, but I stopped.			
Response	Src.	\mathcal{G}_{LL}	\mathcal{G}_{MI}
Why did you stop? was it too dirty? Do you have a job now?	\mathcal{G}	-1.38	1.16
That must have been an exciting job, but why did you quit? I love getting fresh meat from the butcher	\mathcal{G}	-1.62	1.37
Hi! You've watched the movie the chronicles of riddick?	D	-2.55	0.10

Table 8: Example responses and corresponding generator scores $\mathcal{G}_{LL}(c, r)$ (LL score) and $\mathcal{G}_{MI}(c, r)$ (MI score) for model-level G2R. *Src.* column indicates the origin of each response, where \mathcal{G} means it is created by the generation model \mathcal{G} at data-level G2R and D means the response is from original dialogue.

in a situation where only CPU is available for inference.

For calculating CPU latency, we utilized a Ubuntu machine with 40 Intel Xeon Silver 4210 CPU (2.20GHz) and 250GB RAM, and measured latency as the average inference time of 50 response generations after having 3 warm-up steps.

D Human Evaluation Details

We provide additional statistics about the human evaluation result, including 95% confidence interval and p-values for two-tailed t-test between the human evaluation scores of two models in Table 9 and Table 10, respectively. Since the number of annotations was relatively small (200 examples) due to the cost of the human evaluation, the majority of the comparison is not statistically significant ($p < 0.05$). However, we observed that the comparison between *G2R-D vs. Bi-encoder (w/ FAISS)*, *G2R-DM vs. Bi-encoder (w/ FAISS)*, *G2R-DM (MI Score) vs. Bi-encoder (w/ FAISS)* and *G2R-DM vs. Poly-encoder* shows a statistically significant difference in terms of Sum of human evaluation score and the Appropriateness human evaluation score, proving that our G2R methods improve the performance of the retrieval model. Also, note that the trend of the *Sum* human evaluation score within 90M, 2.7B, and 9.4B Blender models is similar to the trend of ACUTE-Eval Engagingness evaluation

result reported in the original paper (Roller et al., 2021), which adds more reliability to our human evaluation result.

E Dataset Details

BST+ dataset is a concatenation of four English dialogue dataset (Blended Skill Talk (Smith et al., 2020), ConvAI2 (Zhang et al., 2018), Empathetic Dialogues (Rashkin et al., 2019) and Wizard of Wikipedia (Dinan et al., 2018)). We use the Blender 9.4B model to augment this dataset as described in the *Data-level G2R* section, and the augmented dataset consists of total 3,070,033 context-response pairs on 274,233 unique contexts. As described in the Experiments section, we release the augmented BST+ dataset in <https://github.com/hyperconnect/g2r>.

F Validation and Test Hits@1/K metrics

For reference, we report the Hits@1/K and the Hits@5/K metrics of our retrieval models measured on the validation and the test split of BST+ in Table 7.

G Details for Model-level G2R Ablation Study

Here, we provide additional details for calculating the MI score in the ablation study for model-level G2R. MI score is calculated with the *MMI-bidi* equation in the original paper (Li et al., 2016), but additionally normalized by the length of response in the same way LL score is normalized:

$$\mathcal{G}_{MI}(c, r) = (\log P_{\mathcal{G}}(r|c) - \log P_{\mathcal{G}}(r)) / |r|. \quad (6)$$

Since calculating the unconditional language probability term $P_{\mathcal{G}}(r)$ in Equation 6 is intractable, we approximate this term by taking the average of the likelihood values of r given dummy input contexts, including ".", "<PAD>" and "<UNK>". This trick enables us to avoid undesirable alternative options for calculating $P_{\mathcal{G}}(r)$ with high computational burden, such as training a separate unconditional language model or calculating an intractable marginal probability $\sum_c P_{\mathcal{G}}(r|c)P(c)$.

Models	Human Evaluation		
	Sum	Appr.	Info.
Blender 90M	2.843±0.091	1.429±0.058	1.414±0.048
Blender 2.7B	2.983±0.091	1.510±0.054	1.473±0.053
Blender 9.4B	2.930±0.092	1.472±0.056	1.458±0.053
Distilled Blender	2.910±0.087	1.474±0.054	1.436±0.051
RetNRef	2.771±0.085	1.404±0.049	1.368±0.053
Bi-encoder	2.597±0.105	1.288±0.060	1.309±0.062
Poly-encoder	2.686±0.094	1.340±0.055	1.346±0.055
Bi-encoder (w/ FAISS)	2.596±0.096	1.259±0.056	1.337±0.055
G2R-D (w/o FAISS)	2.779±0.100	1.380±0.056	1.399±0.059
G2R-D	2.759±0.109	1.398±0.060	1.361±0.064
G2R-DM	2.856±0.098	1.447±0.058	1.410±0.056
G2R-DM (MI Score)	2.806±0.098	1.427±0.059	1.380±0.056
Human Response	2.788±0.103	1.418±0.058	1.369±0.060

Table 9: Human evaluation results of the baseline models and our G2R models with 95% confidence interval.

Model A	Model B	Human Evaluation		
		Sum	Appr.	Info.
G2R-D	Bi-encoder (w/ FAISS)	0.028	0.001	0.575
G2R-DM	Bi-encoder (w/ FAISS)	<0.001	<0.001	0.073
G2R-DM (MI Score)	Bi-encoder (w/ FAISS)	0.003	<0.001	0.291
G2R-D	Poly-encoder	0.311	0.156	0.717
G2R-DM	Poly-encoder	0.014	0.009	0.112
G2R-DM (MI Score)	Poly-encoder	0.080	0.035	0.396
G2R-DM	G2R-D	0.199	0.255	0.266
G2R-DM (MI Score)	G2R-D	0.533	0.507	0.672
Human Response	G2R-D	0.715	0.636	0.860
G2R-DM	G2R-DM (MI Score)	0.481	0.636	0.459
Blender 90M	Human Response	0.427	0.793	0.252
Distilled Blender	Human Response	0.078	0.176	0.097
Blender 2.7B	Human Response	0.006	0.024	0.011
Blender 9.4B	Human Response	0.044	0.194	0.029
G2R-DM	Blender 90M	0.851	0.679	0.907
Human Response	G2R-DM	0.346	0.502	0.335
Distilled Blender	G2R-DM	0.428	0.505	0.499
Blender 2.7B	G2R-DM	0.064	0.116	0.110

Table 10: P-value of the two-tailed t-test between two models on human evaluation results. We boldface the p-values under 0.05. Model A has a better average *Sum* human evaluation score than Model B.

H Data-level and Model-level G2R Examples

Table 8 shows the example responses generated by the data-level G2R, and LL and MI score calculated for each response. Data-level G2R is able to generate high-quality responses that are appropriately related to the input context. Model-level G2R helps the retrieval model to distinguish between low-quality and high-quality responses since both LL and MI scores assign a low value for an inappropriate response. Comparing both scores in this example, the MI score assigns higher values to a relatively more specific response compared to the LL score. In comparison, the LL score assigns higher values to a relatively more general response

that can also be appropriately utilized for a different context.

I Additional Response Examples

We provide additional examples of responses generated by the baseline models and our G2R models in Table 11. Also, the full dialogue context for the case study example in the main paper is provided in 12.

Input Context	
A: I've been doing it as a hobby for about five years. I'm still not great but I keep improving.	
B: That is nice!!! I am usually very punctual. I make it early to every class.	
A: That's a great quality to have. It will help you move up later in life. When you start working that will make a great impact on your boss.	
Model Name	Response
Bi-encoder+FAISS	B: Definitely. Hopefully lead to a promotion for you soon!
Blender 90M	B: That is very true. I will have to keep that in mind when I start my new job.
Distilled Blender	B: I hope so. I need to get my grades up so I can get a good job.
Blender 2.7B	B: I hope so. I want to be the best at what I do. What do you do for a living?
Blender 9.4B	B: I hope so. I want to open my own studio one day. What do you do for a living?
G2R-D	B: Yes it should. I can't wait to get started. I start in august.
G2R-DM	B: I think productivity is important. I am glad I have the opportunity to work towards something.
G2R-DM (MI Score)	B: That is great advice! I like to work hard and not let things bother me.

Input Context	
(Further dialogue history omitted for space)	
A: it is ok I am not in a hurry to be honest	
B: were you in the military	
A: never been in the military how about you ?	
Model Name	Response
Bi-encoder+FAISS	B: Not in the army, but I was in the fbi, g13 classified
Blender 90M	B: I was in the army for a few years before I joined the air force. It is a lot of work
Distilled Blender	B: I was in the navy for a few years. What do you do for a living?
Blender 2.7B	B: I was in the army for 8 years. I was a medic. What did you do?
Blender 9.4B	B: Yes I was in the marine corps for 6 years and now I am in the reserves
G2R-D	B: No I never did but I was thinking about it. Do you like the army? Do u have any pets?
G2R-DM	B: I was in it for 10 years and 2 deployments to iraq and afghanistan
G2R-DM (MI Score)	B: I was in a combat zone but was never deployed

Input Context	
A: I have this cat that loves to cause issues	
B: What sort of issues?	
A: Like last night. She decided that she wanted to jump on the table	
Model Name	Response
Bi-encoder+FAISS	B: Oh no. Pets are good for anger.
Blender 90M	B: Oh no, that's not good. I'm glad you're able to take care of her.
Distilled Blender	B: Oh no! I hope she doesn't do that again. Do you have any other pets?
Blender 2.7B	B: Oh no! I hope she doesn't do that too often. Cats can be pretty destructive.
Blender 9.4B	B: Oh no! I hope she didn't do too much damage. Cats can be so destructive.
G2R-D	B: That's hilarious. Is she a weird cat? I have a cat that will do the same thing if he's not feeling well or doesn't get what he feels like he should.
G2R-DM	B: Oh yea, cats tend to do that!
G2R-DM (MI Score)	B: Oh no! That doesn't really sound like a good time. What is her name? Mine is named einstein.

Table 11: Additional example responses in the BST+ test set example.

Input Context	
B: Hey, jefferson here, I love documentaries	
A: I love dogs her name is barnaby, a german shepherd	
B: Oh that's so cool. I have a dog with my husband in pennsylvania.	
A: Nice to meet you, I work in the usps	
B: I order a lot of books online. I read one every week.	
A: How big are the dogs? Are they grey color? That's my color	
B: I would not know, they are your dogs haha.	
A: Do you order from amazon? My dad is a nascar racer	
B: I do order from amazon, I'm a pa right now.	
A: Are there scorpions in your area? I hate them	
B: I sure hope not, they are scary!	
A: Amazon is a great place to order books	
B: Oh yeah that is a fact.	
Model Name	Response
Bi-encoder (w/ FAISS)	A: Comics is at the top of the list in interest
Blender 90M	A: Do you have any other hobbies that you like to do on the weekends? I like to read
Distilled Blender	A: What else do you like to do in your spare time? I like to go to the beach
Blender 2.7B	A: What kind of books do you buy on amazon? I like to read science fiction and fantasy
Blender 9.4B	A: What kind of books do you like to read? I like romance novels and mystery novels
G2R-D	A: Do you ever order groceries? I love amazon's selection of books and videos
G2R-DM	A: What is your favourite book? Mine is "the cat in the hat" by dr seuss

Table 12: Example responses with full dialogue history in the case study example.