

Transferable Persona-Grounded Dialogues via Grounded Minimal Edits

Chen Henry Wu¹, Yinhe Zheng², Xiaoxi Mao³, Minlie Huang¹

¹ Department of Computer Science and Technology, 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 Samsung Research China - Beijing (SRC-B) ³ Fuxi AI Lab, NetEase Inc., Hangzhou, China
henrychenwu@cmu.edu, yh.zheng@samsung.com,
maoxiaoxi@corp.netease.com, aihuang@tsinghua.edu.cn

Abstract

Grounded dialogue models generate responses that are grounded on certain concepts. Limited by the distribution of grounded dialogue data, models trained on such data face the *transferability* challenges in terms of the data distribution and the type of grounded concepts. To address the challenges, we propose the *grounded minimal editing* framework, which minimally edits existing responses to be grounded on the given concept. Focusing on personas, we propose Grounded Minimal Editor (GME), which learns to edit by disentangling and recombining persona-related and persona-agnostic parts of the response. To evaluate persona-grounded minimal editing, we present the PERSONAMINEDIT dataset, and experimental results show that GME outperforms competitive baselines by a large margin. To evaluate the transferability, we experiment on the test set of BLEND-EDSKILLTALK and show that GME can edit dialogue models' responses to largely improve their persona consistency while preserving the use of knowledge and empathy.¹

1 Introduction

Grounding dialogue agents on external information is important for building engaging conversational AI systems (Huang et al., 2020). Along this track, various datasets and models have been proposed to ground dialogues on personas (Zhang et al., 2018), knowledge (Dinan et al., 2019), emotions (Zhou et al., 2018a), and images (Shuster et al., 2020).

Generally, grounded dialogue modeling trains a dialogue model on a dataset \mathcal{D} that consists of triples (c, r, g) , where c is the dialogue history, r is the response, and g is the grounded concept. The model is generally optimized using maximum likelihood estimate (MLE), i.e.,

$$\arg \max_{\theta} \mathbb{E}_{(c,r,g) \sim \mathcal{D}} \log P_{\theta}(r|c, g). \quad (1)$$

¹Our codes and data are available at <https://github.com/thu-coai/grounded-minimal-edit>.

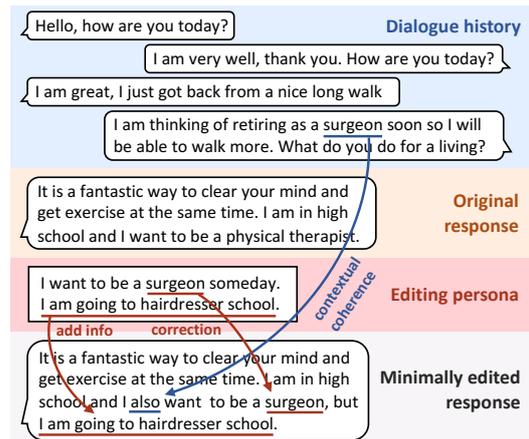


Figure 1: Persona-grounded minimal editing. Edits are shown by arrows, accompanied by the explanations.

Despite its effectiveness, this formulation faces two challenges regarding *transferability*. On one hand, grounded dialogue datasets are usually collected under a guided setting, e.g., annotators are usually encouraged to embed persona (Zhang et al., 2018) or knowledge (Dinan et al., 2019) into responses, which leads to a *distributional gap* between the conversations in a grounded dialogue dataset and natural conversations. As a result, models trained with Eq. (1) may generate unnatural responses and are vulnerable to the distributional shift of the dialogue history. On the other hand, at inference time, models trained with Eq. (1) cannot be grounded on unseen types of concept g' other than g . An example for such *grounding gap* is that a model trained on PERSONACHAT (Zhang et al., 2018) with Eq. (1) cannot be grounded on world knowledge.

To address the above transferability challenges, we propose a *grounded minimal editing* framework for grounded dialogue modeling. Instead of learning a *grounded response generator* as is done in Eq. (1), we propose to learn a *grounded minimal editor* that operates on existing responses. Specifically, suppose we have an original response r^o that is coherent with the dialogue history c but is

not grounded on the concept g . Our goal is to *minimally* edit r^o such that it is grounded on the concept g and coherent with the dialogue history c . Original responses can be generated by dialogue models trained on natural conversation data and grounded on other concepts g' , or even produced by humans; thus, they do not suffer from the *distributional gap* and *grounding gap*. Moreover, *minimal* editing guarantees that the distribution of the edited responses is similar to that of the original responses, which do not suffer from the two gaps. Note that collecting paired responses before and after editing is resource-consuming; thus, our goal is to learn the editing *without* paired data.

In this paper, we explore *persona-grounded minimal editing*, as demonstrated in Figure 1. We propose Grounded Minimal Editor (GME), which is trained on persona-grounded dialogue data. Specifically, response templates are sampled by corrupting persona-related spans and sentences based on gradient-based attribution and word overlap. By denoising the templates, GME disentangles and recombines persona-related and persona-agnostic expressions. Since the personas of original responses are not observed at inference, we train a classifier for template generation at inference.

Two research questions are investigated in this paper: Q1) *Is the proposed GME model effective for grounded minimal editing?* Q2) *Does our framework address the transferability challenges (more specifically, the distributional gap and the grounding gap)?* For Q1, we build PERSONAMINEDIT, a new dataset derived from PERSONACHAT with multiple human references for the edited response. Automatic and human evaluations show that GME outperforms competitive baselines and has the most similar behavior to humans references. For Q2, we evaluate GME on the test set of BLENDEDSKILLTALK (Smith et al., 2020), whose data distribution and grounded concepts are different from PERSONACHAT, which requires GME to be transferable. We observe that GME improves the persona consistency of responses generated by pretrained Blender-90M models (Roller et al., 2020), while preserving the use of knowledge and empathy. Results also show that GME-edited responses largely outperforms TransferTransfo (Wolf et al., 2019), which is trained in the canonical way as in Eq. (1). Our contributions include:

- We propose a framework named *grounded minimal editing* to address the transferability

challenges of grounded dialogue modeling.

- We propose Grounded Minimal Editor (GME) and present the PERSONAMINEDIT dataset to evaluate GME’s effectiveness for persona-grounded minimal editing.
- Experimental results show that GME largely outperforms strong baselines on the PERSONAMINEDIT dataset. GME is also transferable to edit other models’ outputs and improve the persona consistency while preserving their use of knowledge and empathy.

2 Related Work

Recent work leveraged grounded information in dialogue agents to chat engagingly, e.g., using knowledge (Zhou et al., 2018b), emotions (Zhou et al., 2018a), personas (Zhang et al., 2018), and images (Shuster et al., 2020). For persona grounding (Li et al., 2016; Zhang et al., 2018), transfer learning methods (Zhang et al., 2019; Wolf et al., 2019; Golovanov et al., 2019) and latent variable models (Song et al., 2019; Chan et al., 2019) have shown promising results. Further, the persona consistency issue (Kim et al., 2020; Nie et al., 2020) and persona-augmented empathetic agents (Zhong et al., 2020) have also been explored. As discussed in Section 1, existing methods generally adopt the MLE objective in Eq. (1) and suffer from two transferability challenges, i.e., the distributional gap and the grounding gap, which are addressed by the proposed *grounded minimal editing* framework.

The idea of *editing* existing responses has been explored, e.g., the deliberation network (Xia et al., 2017), two-pass response generation (Song et al., 2020), and retrieval-augmented dialogue modeling (Weston et al., 2018; Pandey et al., 2018; Wu et al., 2019b; Gu et al., 2019; Cai et al., 2019). This paper is essentially different from these works from two perspectives. 1) Regarding the formulation, we emphasize *minimal* editing, while previous works do not. As analyzed in Section 1, *minimal* editing is an important component to address the transferability challenges; 2) Regarding the training algorithm, previous works derive templates from self-generated or retrieved texts, while our model derives templates from the observed responses.

Our work is also related to controlled text editing without parallel data, e.g., unsupervised text style transfer (Shen et al., 2017; Li et al., 2018; Rao

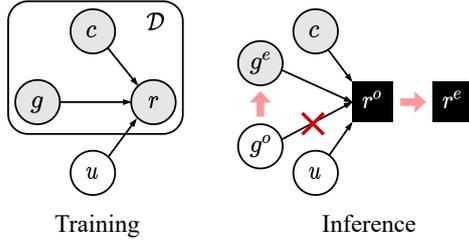


Figure 2: Graphical formulation of grounded minimal editing. Observed variables are shown in grey, and unobserved variables are shown in white (c : dialogue history, g : grounding, r : response, \mathcal{D} : training data, u : unobserved variables). At inference, the editing $r^o \rightarrow r^e$ is based on $g^o \rightarrow g^e$, while the unobserved variables u remain unchanged. Note that g^o is also not observed.

and Tetreault, 2018; Lample et al., 2019), semi-supervised contextual text style transfer (Cheng et al., 2020), syntax-controlled paraphrasing (Bao et al., 2019), contrastive model explanation (Ross et al., 2020), counterfactual story generation (Qin et al., 2019, 2020), and sentence-level editing for empathetic dialogues (Sharma et al., 2021). Some of these studies also utilize masked templates (Li et al., 2018; Wu et al., 2019a; Sudhakar et al., 2019; Malmi et al., 2020; Ross et al., 2020). However, these previous works only focus on categorical conditions in a small label space, while the personas in our study are embedded in much larger spaces. In the large persona space, the persona sentences at test time are never seen during training. Further, when generating masked templates, the personas of the original responses are unobserved in our study.

3 Formulation

We provide a formulation of the proposed framework. Grounded dialogue modeling uses a dataset \mathcal{D} that consists of triples (c, r, g) , where c , r , and g are the dialogue history, the response, and the grounded concept, which are shown in grey in the left part of Figure 2. To formulate the term “minimal”, we need to add *unobserved variables* into the graphical model, denoted as u in Figure 2, which cover *all* unobserved variables. The graph states that $r = f(c, g, u)$. As shown in the right part of Figure 2, we observe (c, r^o, g^e) at inference time, where r^o and g^e stand for the original response and the grounded concept for editing. The graph states that the original response $r^o = f(c, g^o, u)$, where g^o represents the concept the original response is grounded on, and that both g^o and u are unobserved. The edited response is defined as $r^e = f(c, g^e, u)$,

which replaces g^o as g^e , and keeps c and u intact. Our formulation follows the idea of counterfactual reasoning (Peters et al., 2017), and it guarantees that 1) the content irrelevant to the grounded concept is preserved, and that 2) the edited response is coherent with the dialogue history. Since it is costly to collect paired (r^o, r^e) for training, the grounded minimal editor should be trained on the grounded dialogue data $(c, r, g) \sim \mathcal{D}$ as in Eq. (1).

As the first attempt toward the proposed framework, we focus on *persona-grounded* minimal editing in the experiments. Thus, in the remaining part of this paper, we set the grounded concept g , g^o , g^e as the persona p , p^o , p^e .

4 Our Approach

4.1 Overview

We propose Grounded Minimal Editor (GME), a pipeline model for grounded minimal editing. At inference, GME first creates a *response template* t by masking persona-related spans in the original response r^o and then *recombines* the template t , the persona p^e , and the dialogue history c into an edited response r^e . We design the template to approximate the unobserved variables u in Section 3, which distinguishes GME from previous retrieval-based dialogue models. With some abuse of notation, we use t to denote the template for both training and inference. During training, two modules are learned: 1) a generator used for the recombination described above and 2) a mask classifier that helps create the response template at inference. Note that GME can also be applied to other ground concepts besides personas. The full process is presented in Algorithm 1.

4.2 Recombination Module

The recombination module learns to recombine the response template, the persona, and the dialogue history as the edited response. During training, we create templates from the training responses, as detailed below.

Span mask The span mask serves as the placeholder of persona-related spans. For each response-persona pair, we define three sets of tokens: GRADIENT, OVERLAP, and STOPWORDS. GRADIENT contains persona-related tokens that are determined using *gradient-based attribution* (Simonyan et al., 2014). We pretrain a response-to-persona model and compute the L_2 norm of the gradient of the

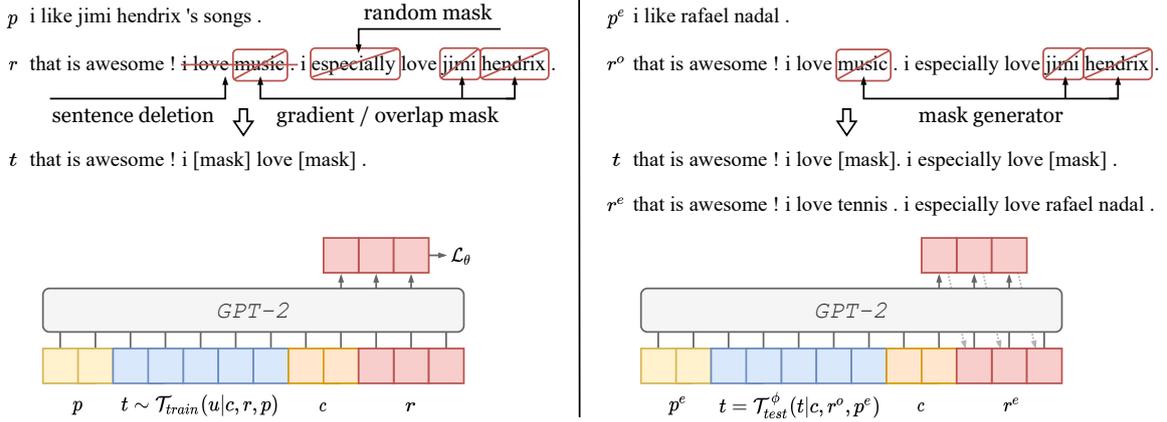


Figure 3: Left: training example and input format. Right: inference example and input format. For readability, the same response sample is used. The context, position embeddings, and token type embeddings are omitted here.

Algorithm 1 Training and inference of GME

```

// Training
1: repeat sample  $(c, r, p) \sim \mathcal{D}$ 
2: Sample  $t \sim \mathcal{T}_{train}(t|r, p)$ 
3: Optimize  $\mathcal{L}_\theta$  in Eq. (2) and  $\mathcal{L}_\phi$  in Eq. (3)
4: until convergence
// Inference
Input:  $(c, r^o, p^e)$ 
5: Infer  $t = \mathcal{T}_{test}^\phi(t|c, r^o, p^e)$ 
6: Edited response  $r^e \sim P_\theta(r^e|c, t, p^e)$ 

```

persona’s cross-entropy loss w.r.t. each response token’s embeddings. A token is placed into the GRADIENT set if the L_2 norm is greater than $\delta = 3$. OVERLAP contains response tokens whose lemma overlaps with in the lemmas of the persona tokens, which are likely to be related to the persona. STOPWORDS contains stopwords and punctuation marks specified by NLTK (Bird, 2006). We mask a token if it is in GRADIENT or OVERLAP but not in STOPWORDS. We call sentences with masks after this step as *persona-related sentences*. For each persona-related sentence, we further mask 15% of its tokens to improve the robustness. Since the number of tokens varies at the same syntactic position, we merge consecutive masks so that all masks are at the span level.

Sentence deletion The above span mask is effective for correcting persona *contradictions* in the original response. However, span mask cannot handle the situation where we want to add new persona information into the response (examples are given in Figure 1 and Appendix E). To model this pattern, we randomly delete persona-related sentences. Suppose we have l persona-related sentences in the response, the number to keep $0 \leq n \leq l-1$ follows

$P(n) \propto \exp(-n/\tau)$, where τ is a hyperparameter. By infilling persona-related sentences, the model learns to merge persona into the response.

An example of the training template is shown in Figure 3. During training, the recombination modules P_θ is optimized by

$$\mathcal{L}_\theta = -\mathbb{E}_{(c,r,p) \sim \mathcal{D}} [\mathbb{E}_{t \sim \mathcal{T}(t|r,p)} \log P_\theta(r|c, t, p)]. \quad (2)$$

where $\mathcal{T}(t|r, p)$ denotes the distribution of the template as detailed above. As shown in Figure 3, we use GPT-2 as the backbone to parameterize P_θ , which is tackled as a language modeling task by concatenating input texts. We apply label smoothing ($\epsilon = 0.1$), and we use greedy decoding at inference. Token type embeddings are used to distinguish each type of text and each speaker.

4.3 Mask Generator

Since the persona of the original response before editing, i.e., p^o , is unobserved at inference, we train a mask generator P_ϕ to predict if a token r_i should be masked. The objective for the mask generator is

$$\mathcal{L}_\phi = -\mathbb{E}_{(c,r,p) \sim \mathcal{D}} \sum_{i=1}^{|r|} \frac{1}{f_i} \log P_\phi(m_i|c, r) \quad (3)$$

where $m_i = 1$ if r_i is in GRADIENT or OVERLAP but not in STOPWORDS, and $m_i = 0$ otherwise. f_i is the corpus-level frequency of m_i , which is used to balance the number of positive samples and negative samples. At inference, we mask a word if 1) P_ϕ labels it as *masked* with a confidence greater than ϵ ($\epsilon = 0.5$ in the main experiment, $\epsilon = 0.75$

in the transferability experiment) and meanwhile 2) it does not appear in the persona p^e or the dialogue history c . We merge consecutive masks to get span masks. This process is denoted as $\mathcal{T}_{test}^\phi(t|c, r^o, p^e)$ in Algorithm 1.

5 Evaluation Data: PERSONAMINEDIT

5.1 Data Collection

We present a new dataset PERSONAMINEDIT to evaluate persona-grounded minimal editing. Validation and test data are collected in two steps:

Editing persona selection We first construct inference samples (c, r^o, p^e) , where the dialogue history c and original response r^o are from PERSONACHAT, and we select the editing persona p^e based on two criteria: 1) *editing difficulty* and 2) *conversation consistency*. We bias our data to the hard cases that require correction of *persona contradictions*. Specifically, we use the heuristics provided by Welleck et al. (2019) to select personas that are contradictory to the original response. To ensure conversation consistency, we filter out personas that are contradictory to the speaker’s responses in the dialogue history. Finally, we also ensure that the persona sentences within each persona are not contradictory to each other.

Response editing For each constructed triple (c, r^o, p^e) , we collect references for the edited responses r^e on Amazon Mechanical Turk. Specifically, r^e should satisfy three requirements: 1) consistency with the editing persona p^e , 2) minimal editing, and 3) coherence with the dialogue history c . We reject annotations that do not add words to the original response. *Three* human references are collected for each triple, and duplicate references are re-annotated. The inter-annotator BLEU (i.e., the BLEU of each reference given the other two references) is 73.8 on the validation set and 71.4 on the test set. The annotation instructions we used are detailed in Appendix A.

Training data in PERSONAMINEDIT is derived from the training data of PERSONACHAT, and personas are aligned with responses following Welleck et al. (2019). We also *remove training samples whose persona appears in the editing personas in the validation and test data* to ensure that the *persona does not leak* from training to testing.

5.2 Data Statistics

After removing training samples whose persona appears in the editing personas in the validation and

	<i>add</i>	<i>rm</i>	ΔL	$d(r^e, r^o)$	$d(r^e, p^e)$
Valid	5.1	2.7	+ 3.2	7.0	11.0
Test	4.9	2.6	+ 2.9	6.7	10.8

Table 1: Data analysis. Notations follow Section 5.2.

test splits, our training data has 119,078 samples. The validation split has 1,384 samples (1,266 with one sentence in the editing persona, 118 with two). The test split also has 1,384 samples (1,269 with one sentence in the editing persona, 115 with two).

We study the *behavior* of human references to understand the human intuition of minimal editing. In Table 1, we report the number of words added (*add*) and removed (*rm*), and the length difference (ΔL) between the edited and original responses. We also report the minimum edit distance (MED) between the edited and original responses ($d(r^e, r^o)$), and that between the edited response and the editing persona ($d(r^e, p^e)$). We observe that the edited responses are generally local modifications of the original responses. On average, the edited responses are longer than the original ones, which can be explained by the observation that human sometimes add persona information into the response when no persona contradiction exists.

6 Experiment on PERSONAMINEDIT

We use PERSONAMINEDIT to evaluate persona-grounded minimal editing (Q1 in Section 1).

6.1 Baselines

We modify state-of-the-art models for unsupervised text style transfer and counterfactual story generation as the baselines for grounded minimal editing.

No edit This baseline does not make any edits to the original response.

UNMT Lample et al. (2019); He et al. (2020) adopted the unsupervised neural machine translation (UNMT) model (Lample et al., 2018) for unsupervised text style transfer. For our task, we replace the *style condition* with *persona condition*, and use a word dropout rate $p_{wd} \in \{0.1, 0.5\}$.

CycleGAN Luo et al. (2019); Dai et al. (2019) adopted CycleGAN (Zhu et al., 2017) for unsupervised text style transfer. For our task, we replace the *style classifier* with a *response-to-persona model*. We use Gumbel-softmax straight through gradient estimator (Jang et al., 2017) for optimization.

DeLorean-FT DeLorean (Qin et al., 2020) iteratively modifying GPT-2’s logits via gradients from

a content preserving loss. For our task, we replace GPT-2 with TransferTransfo (Wolf et al., 2019) and set the mixture rate $\gamma_{mix} \in \{0.75, 0.80, 0.85\}$, where larger (smaller) γ_{mix} is biased towards persona consistency (minimality of editing).

We observe that CycleGAN is sensitive to hyperparameters and unstable to train, probably due to the biased gradient estimation given the large persona space. Thus, we do not include other methods that require gradient backpropagation from classifiers (Zhou et al., 2020; Madaan et al., 2020).

6.2 Automatic Evaluation

For automatic evaluation, we run each experiment with five random seeds. More details are presented in Appendix C.

BLEU We compute BLEU-4 score (Papineni et al., 2002) based on the collected multiple human references, using the Moses script *multi-bleu.perl*. From Table 2 and Table 5, we observe that higher BLEU indicates the less editing.

P-Score We define P-Score to evaluate the persona consistency. Specifically, we finetune a BERT model on the DNLI dataset (Welleck et al., 2019) to predict the relation $C(r, p_j)$ (*entailment*, *neutral*, or *contradiction*) of a response r and a persona sentence p_j .² We then map *entailment*, *neutral*, and *contradiction* to +0.5, 0, and -0.5 and define the P-Score of a sample as

$$\text{P-Score} = \sum_j \text{map}[C(r^e, p_j^e)] \quad (4)$$

where r^e is the edited response and p_j^e is a persona sentence in p^e . We finally report the P-Score averaged over all samples.

Average We observe that BLEU and P-Score show a trade-off between minimal editing and persona consistency. We report their arithmetic mean as the overall performance since BLEU and P-Score have similar scales and variances.

Table 2 shows that CycleGAN and UNMT have high BLEU but negative P-Scores. Figure 4 shows that most of their outputs are contradictory to the editing personas, indicating that their edits are not focused on persona-related expressions. These results show that methods designed for binary style labels are not effective for persona-grounded minimal editing, where the persona space is much larger than the label space. Larger γ_{mix} for DeLorean-FT

²We use the classifier provided by (Madotto et al., 2019), which has 92.57% accuracy on the DNLI verified test set.

	BLEU	P-Score	Average
No edit	76.4 (0.0)	- 30.5 (0.0)	23.0 (0.0)
UNMT			
- $p_{wd} = 0.1$	74.2 (0.2)	- 30.2 (0.3)	22.0 (0.2)
- $p_{wd} = 0.5$	69.0 (0.2)	- 27.9 (0.7)	20.6 (0.4)
CycleGAN	74.4 (0.8)	- 28.3 (1.6)	23.0 (0.7)
DeLorean-FT			
- $\gamma_{mix} = 0.75$	39.8 (2.2)	+ 26.4 (2.8)	33.1 (0.6)
- $\gamma_{mix} = 0.80$	34.5 (0.7)	+ 32.6 (1.6)	33.5 (0.6)
- $\gamma_{mix} = 0.85$	32.0 (0.8)	+ 36.5 (1.0)	34.2 (0.7)
GME (ours)	60.3 (1.8)	+ 29.9 (2.2)	45.1 (0.5)

Table 2: Automatic evaluation. We report the average of 5 random seeds, and standard deviations are shown in parenthesis. Details of P-Score are in Figure 4.

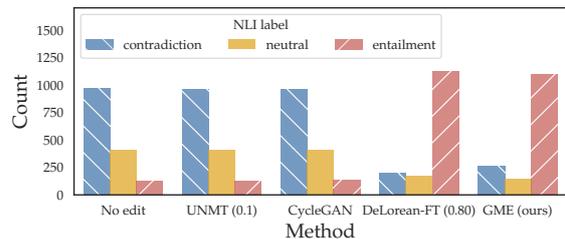


Figure 4: Distribution of classes in P-Score: contradiction (blue backslash), neutral (yellow), and entailment (red slash). Hyperparameters are in parenthesis.

lead to lower BLEU and higher P-Score, showing that larger (smaller) γ_{mix} is biased towards persona consistency (minimality of editing). However, results show that the overall performance cannot be improved by hyperparameter tuning.

GME achieves a 31.9% relative improvement on the Average score over the best performing baseline (from 34.2 to 45.1). Figure 4 shows that most of GME’s outputs entail the given personas. Table 4 shows the results for 1) removing *dialogue histories* from the data and 2) removing *sentence deletion* from GME. We observe that the dialogue history only has a slight contribution, showing that the response template contains an adequate amount of information of the original response. Sentence deletion contributes largely to the performance, especially for the persona consistency.

6.3 Human Evaluation

We randomly sample 150 test samples for human evaluation. Given two edited responses A and B , three annotators are hired to vote *prefer A*, *none*, or *prefer B*. We instruct annotators vote *none* if neither A nor B satisfies both *minimal editing* and *persona consistency*. See detailed guidelines in Appendix B and supplementary materials.

Baseline	<i>prefer baseline</i>	<i>none</i>	<i>prefer GME</i>
UNMT (0.1)	0.0 %	59.3 %	40.7 %
CycleGAN	0.7 %	59.1 %	40.2 %
DeLorean-FT (0.85)	8.4 %	52.7 %	38.9 %

Table 3: Human evaluation. Free-marginal κ for each row is 0.66, 0.66, and 0.51 (*substantial, substantial, and moderate agreement*).

	BLEU	P-Score	Average
Full	60.3 (1.8)	+ 29.9 (2.2)	45.1 (0.5)
w/o history	60.2 (1.0)	+ 29.3 (2.0)	44.8 (0.8)
w/o sent. del.	64.2 (0.3)	+ 11.0 (0.3)	37.6 (0.2)

Table 4: Ablation studies. Notations follow Table 2.

Table 3 shows that human annotators generally prefer GME to the baselines. The free-marginal κ for each row is 0.66, 0.66, and 0.51 (*substantial, substantial, and moderate agreement*). The strongest baseline DeLorean-FT is only preferred in 8.4% cases. We observe that in most cases where DeLorean-FT wins, the original response is syntactically similar to the persona.

6.4 Behavioral Analysis

Using the metrics defined in Section 5.2, we provide a behavioral analysis of the models. Results are shown in Table 5. CycleGAN and UNMT have small *add*, *rm*, and $d(r^e, r^o)$, which shows that they make little changes to the original response. For DeLorean-FT, larger mixture rates γ_{mix} have larger *add*, *rm*, and $d(r^e, r^o)$, which is consistent with the observation in Section 6.2. The large $d(r^e, r^o)$ of DeLorean-FT also shows that this model behaves poorly at making minimal editing. GME has the most similar behavior with human references. Based on the observations in Section 6.2-6.4, we conclude that GME is effective in making minimal edits that are targeted at persona-related expressions. By checking the outputs, we observe that GME and human references add persona information into the response in some cases, which may explain why GME and human references have positive ΔL (i.e., their predictions are longer than the original responses).

7 Transferability Experiment

We evaluate the transferability of GME to minimally edit existing responses on the test split of BLENDEDSKILLTALK (Smith et al., 2020). We also evaluate whether grounded minimal editing ad-

	<i>add</i>	<i>rm</i>	ΔL	$d(r^e, r^o)$	$d(r^e, p^e)$
No edit	0.0	0.0	0.0	0.0	12.0
UNMT (0.1)	0.2	0.1	+ 0.1	0.2	12.1
UNMT (0.3)	0.6	0.4	+ 0.3	1.0	12.2
CycleGAN	0.1	0.1	+ 0.1	0.2	12.1
DeLorean-FT					
- $\gamma_{mix} = 0.75$	5.7	7.9	- 1.7	9.8	7.5
- $\gamma_{mix} = 0.80$	6.3	8.7	- 1.8	10.9	7.2
- $\gamma_{mix} = 0.85$	6.9	9.2	- 1.7	11.7	7.0
GME (ours)	4.0	2.5	+ 3.6	6.7	13.0
Human references	4.9	2.6	+ 2.9	6.7	10.8

Table 5: Behavioral statistics. Results that are **the closest to human references** are shown in bold.

dresses the transferability challenges of grounded dialogue modeling (Q2 in Section 1).

7.1 Experimental Setup

In BLENDEDSKILLTALK, each dialogue session is grounded on *two persona sentences* and an optional *knowledge topic*, and the distribution of responses is biased towards the mixture of *displaying persona, using knowledge, and being empathetic*. Two types of existing responses are considered:

- Responses generated by a *persona-agnostic* Blender-90M (Roller et al., 2020), which is trained on BLENDEDSKILLTALK in which the persona sentences are removed.
- Responses generated by the original *persona-grounded* Blender-90M.

We compare the above two Blender-90M variants and GME-edited responses with TransferTransfo (Wolf et al., 2019), a pretrained dialogue model finetuned on PERSONACHAT. Note that GME is *not* finetuned on BLENDEDSKILLTALK. Also, conversations in PERSONACHAT, on which GME and TransferTransfo are trained, barely display knowledge and empathy.

7.2 Automatic Evaluation

We report BLEU and F1 (Miller et al., 2017) computed with the human references. For persona consistency, we report the P-Score defined in Section 6.2. To evaluate fluency, we report the word-level NLL evaluated by GPT-2 (Radford et al., 2018). The automatic evaluation uses the full 5482 test samples of BLENDEDSKILLTALK.

Table 6 shows that P-Scores are largely improved after GME editing (from 9.2 to 33.0, and from 0.8 to 29.4). BLEU, F1, and NLL remain comparable

	Automatic evaluation				Human evaluation			
	BLEU	F1	P-Score	NLL ↓	Knowledge	Empathy	Persona	Grammaticality
TransferTransfo	2.31	13.96	+ 67.4	4.54	0.0 %	4.0 %	82.7 %	90.7 %
Blender-90M	3.23	19.22	+ 9.2	3.73	3.3 %	29.3 %	20.7 %	96.3 %
+ edited by GME	3.10	19.02	+ 33.0	3.82	3.0 %	29.0 %	66.3 %	88.3 %
Blender-90M w/o persona	2.98	18.91	+ 0.8	3.63	8.3 %	28.7 %	1.3 %	97.0 %
+ edited by GME	2.84	18.87	+ 29.4	3.78	7.7 %	28.0 %	56.7 %	91.7 %

Table 6: Automatic and human evaluation for the transferability to the test set of BLENDED SKILL TALK. NLL is computed using GPT-2. Free-marginal κ for knowledge, empathy, persona, and grammaticality is 0.92, 0.70, 0.85, and 0.78 (*almost perfect, substantial, almost perfect, and substantial agreement*).

to those before editing. Although TransferTransfo has the highest persona consistency, it has much poorer BLEU, F1, and NLL than GME. These results show that grounded minimal editing addresses the transferability issue faced by TransferTransfo.

7.3 Human Evaluation

We randomly sample 100 test samples for human evaluation. Three annotators evaluate if a response shows *knowledge*, *empathy*, *persona consistency*, and *grammaticality*. See detailed guidelines in Appendix B and supplementary materials.

Results are shown in Table 6. Free-marginal κ for knowledge, empathy, persona, and grammaticality is 0.92, 0.70, 0.85, and 0.78 (*almost perfect, substantial, almost perfect, and substantial agreement*), respectively. Results show that persona consistency is largely improved after GME editing, while the use of knowledge and empathy remain comparable to those before editing. TransferTransfo has the highest persona consistency, but it has much lower knowledge and empathy than the responses edited by GME. For example, only 4.0% of TransferTransfo’s responses show empathy, while the ratios are 29.0% and 28.0% for the GME-edited responses. We also notice a slight grammaticality drop after GME editing. However, the GME edited responses still achieve competitive or higher grammaticality scores comparing to Transfertransfo. In practice, the grammaticality scores can be easily improved using re-ranking approaches. In summary, GME largely improves the persona consistency of existing responses while preserving their use of knowledge and empathy, which addresses the transferability challenges faced by grounded dialogue models trained on PERSONACHAT, e.g., TransferTransfo.

8 Discussion

As mentioned in Section 2, the term “*minimal*” distinguishes our work from two-pass generation (Xia et al., 2017) and retrieval-augmented dialogue models (Weston et al., 2018; Cai et al., 2019). Generally, their objective can be formulated as $P(r|c, r')$ where r' is a response either generated by the model itself or retrieved from the dataset. However, these works do not require r and r' to be a minimal editing pair. By contrast, we formulate r^e and r^o to be a minimal editing pair. To encourage *minimal* editing, we construct response templates from the observed responses themselves, while these works derive templates from the r' defined above.

GME itself is also trained on a grounded dialogue dataset that has biased distribution. Thus, as we mentioned at the beginning of Section 7, we also need to evaluate the *transferability* of GME. Section 7 shows that GME editing only slightly changes the distribution of the responses generated by the Blender-90M variants, while the distribution of TransferTransfo’s responses is further away from the human references. This observation suggests that *minimally editing* out-of-domain responses is easier than *generating* them.

While we focus on the persona, other types of grounding, e.g., *knowledge* and *image*, remain to be explored. Many of GME’s failure cases (see Appendix E) contain grammatical errors or fail to correct contradictions, which could be addressed by improving the quality of response templates or incorporating stronger language model priors.

9 Conclusions

We propose a framework named *grounded minimal editing* to address the *transferability* challenges of grounded dialogue modeling, which include the distributional gap and the grounding gap. Our Grounded Minimal Editor (GME) model achieves

minimal editing by disentangling and recombining persona-related and persona-irrelevant expressions. For evaluation, we present the PERSONAMINEDIT dataset with multiple human references. Experimental results show the effectiveness of GME for persona-grounded minimal editing. GME is also transferable to edit responses generated by pre-trained dialogue models and improve their persona consistency while preserving their use of knowledge and empathy.

Acknowledgements

This work was supported by the National Science Foundation for Distinguished Young Scholars (with No. 62125604) and 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.

References

- Yu Bao, Hao Zhou, Shujian Huang, Lei Li, Lili Mou, Olga Vechtomova, Xin-yu Dai, and Jiajun Chen. 2019. Generating sentences from disentangled syntactic and semantic spaces. In *Proceedings of ACL 2019*, pages 6008–6019, Florence, Italy.
- Steven Bird. 2006. NLTK: the natural language toolkit. In *Proceedings of ACL 2006, Sydney, Australia, 17-21 July 2006*.
- Deng Cai, Yan Wang, Wei Bi, Zhaopeng Tu, Xiaojiang Liu, Wai Lam, and Shuming Shi. 2019. Skeleton-to-response: Dialogue generation guided by retrieval memory. In *Proceedings of NAACL-HLT 2019, Minneapolis, MN, USA, June 2-7, 2019, Volume 1 (Long and Short Papers)*, pages 1219–1228.
- Zhangming Chan, Juntao Li, Xiaopeng Yang, Xiuying Chen, Wenpeng Hu, Dongyan Zhao, and Rui Yan. 2019. Modeling personalization in continuous space for response generation via augmented wasserstein autoencoders. In *Proceedings of EMNLP-IJCNLP 2019, Hong Kong, China, November 3-7, 2019*, pages 1931–1940.
- Yu Cheng, Zhe Gan, Yizhe Zhang, Oussama Elachqar, Dianqi Li, and Jingjing Liu. 2020. Contextual text style transfer. *CoRR*.
- Ning Dai, Jianze Liang, Xipeng Qiu, and Xuanjing Huang. 2019. Style transformer: Unpaired text style transfer without disentangled latent representation. In *Proceedings of ACL 2019, Florence, Italy, July 28- August 2, 2019, Volume 1: Long Papers*, pages 5997–6007.
- Emily Dinan, Stephen Roller, Kurt Shuster, Angela Fan, Michael Auli, and Jason Weston. 2019. Wizard of wikipedia: Knowledge-powered conversational agents. In *ICLR 2019, New Orleans, LA, USA, May 6-9, 2019*.
- Sergey Golovanov, Rauf Kurbanov, Sergey I. Nikolenko, Kyryl Truskovskiy, Alexander Tselousov, and Thomas Wolf. 2019. Large-scale transfer learning for natural language generation. In *Proceedings of ACL 2019, Florence, Italy, July 28- August 2, 2019, Volume 1: Long Papers*, pages 6053–6058.
- Jia-Chen Gu, Zhen-Hua Ling, Xiaodan Zhu, and Quan Liu. 2019. Dually interactive matching network for personalized response selection in retrieval-based chatbots. In *Proceedings of EMNLP-IJCNLP 2019, Hong Kong, China, November 3-7, 2019*, pages 1845–1854.
- Junxian He, Xinyi Wang, Graham Neubig, and Taylor Berg-Kirkpatrick. 2020. A probabilistic formulation of unsupervised text style transfer. In *ICLR 2020, Addis Ababa, Ethiopia, April 26-30, 2020*.
- Minlie Huang, Xiaoyan Zhu, and Jianfeng Gao. 2020. Challenges in building intelligent open-domain dialog systems. *ACM Trans. Inf. Syst.*, 38(3):21:1–21:32.
- Eric Jang, Shixiang Gu, and Ben Poole. 2017. Categorical reparameterization with gumbel-softmax. In *ICLR 2017, Toulon, France, April 24-26, 2017*.
- Hyunwoo Kim, Byeongchang Kim, and Gunhee Kim. 2020. Will I sound like me? improving persona consistency in dialogues through pragmatic self-consciousness. In *Proceedings of EMNLP 2020*, pages 904–916, Online.
- Diederik P. Kingma and Jimmy Ba. 2015. Adam: A method for stochastic optimization. In *ICLR 2015, San Diego, CA, USA, May 7-9, 2015*.
- Guillaume Lample, Myle Ott, Alexis Conneau, Ludovic Denoyer, and Marc’Aurelio Ranzato. 2018. Phrase-based & neural unsupervised machine translation. In *Proceedings of EMNLP 2018, Brussels, Belgium, October 31 - November 4, 2018*, pages 5039–5049.
- Guillaume Lample, Sandeep Subramanian, Eric Michael Smith, Ludovic Denoyer, Marc’Aurelio Ranzato, and Y-Lan Boureau. 2019. Multiple-attribute text rewriting. In *ICLR 2019, New Orleans, LA, USA, May 6-9, 2019*.
- Jiwei Li, Michel Galley, Chris Brockett, Georgios P. Spithourakis, Jianfeng Gao, and William B. Dolan. 2016. A persona-based neural conversation model. In *Proceedings of ACL 2016, August 7-12, 2016, Berlin, Germany, Volume 1: Long Papers*.

- Juncen Li, Robin Jia, He He, and Percy Liang. 2018. Delete, retrieve, generate: a simple approach to sentiment and style transfer. In *Proceedings of NAACL-HLT 2018, New Orleans, Louisiana, USA, June 1-6, 2018, Volume 1 (Long Papers)*, pages 1865–1874.
- Fuli Luo, Peng Li, Jie Zhou, Pengcheng Yang, Baobao Chang, Xu Sun, and Zhifang Sui. 2019. A dual reinforcement learning framework for unsupervised text style transfer. In *Proceedings of IJCAI 2019, Macao, China, August 10-16, 2019*, pages 5116–5122.
- Nishtha Madaan, Inkit Padhi, Naveen Panwar, and Dip-tikalyan Saha. 2020. Generate your counterfactuals: Towards controlled counterfactual generation for text. *arXiv preprint arXiv:2012.04698*.
- Andrea Madotto, Zhaojiang Lin, Chien-Sheng Wu, and Pascale Fung. 2019. Personalizing dialogue agents via meta-learning. In *Proceedings of ACL 2019, Florence, Italy, July 28- August 2, 2019, Volume 1: Long Papers*, pages 5454–5459.
- Eric Malmi, Aliaksei Severyn, and Sascha Rothe. 2020. Unsupervised text style transfer with padded masked language models. In *Proceedings of EMNLP 2020, Online, November 16-20, 2020*, pages 8671–8680.
- Alexander H. Miller, Will Feng, Dhruv Batra, Antoine Bordes, Adam Fisch, Jiasen Lu, Devi Parikh, and Jason Weston. 2017. Parlai: A dialog research software platform. In *Proceedings of EMNLP 2017, Copenhagen, Denmark, September 9-11, 2017 - System Demonstrations*, pages 79–84.
- Yixin Nie, Mary Williamson, Mohit Bansal, Douwe Kiela, and Jason Weston. 2020. I like fish, especially dolphins: Addressing contradictions in dialogue modelling. *arXiv preprint arXiv:2012.13391*.
- Gaurav Pandey, Danish Contractor, Vineet Kumar, and Sachindra Joshi. 2018. Exemplar encoder-decoder for neural conversation generation. In *Proceedings of ACL 2018, Melbourne, Australia, July 15-20, 2018, Volume 1: Long Papers*, pages 1329–1338.
- Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. 2002. Bleu: a method for automatic evaluation of machine translation. In *Proceedings of ACL 2002, July 6-12, 2002, Philadelphia, PA, USA.*, pages 311–318.
- Jonas Peters, Dominik Janzing, and Bernhard Schölkopf. 2017. *Elements of Causal Inference: Foundations and Learning Algorithms*. Adaptive Computation and Machine Learning. MIT Press, Cambridge, MA.
- Lianhui Qin, Antoine Bosselut, Ari Holtzman, Chandra Bhagavatula, Elizabeth Clark, and Yejin Choi. 2019. Counterfactual story reasoning and generation. In *Proceedings of EMNLP-IJCNLP 2019, Hong Kong, China, November 3-7, 2019*, pages 5042–5052.
- Lianhui Qin, Vered Shwartz, Peter West, Chandra Bhagavatula, Jena D. Hwang, Ronan Le Bras, Antoine Bosselut, and Yejin Choi. 2020. Back to the future: Unsupervised backprop-based decoding for counterfactual and abductive commonsense reasoning. In *Proceedings of EMNLP 2020, Online, November 16-20, 2020*, pages 794–805.
- Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, and Ilya Sutskever. 2018. Language models are unsupervised multitask learners.
- Sudha Rao and Joel R. Tetreault. 2018. Dear sir or madam, may I introduce the GYAFC dataset: Corpus, benchmarks and metrics for formality style transfer. In *Proceedings NAACL-HLT 2018, New Orleans, Louisiana, USA, June 1-6, 2018, Volume 1 (Long Papers)*, pages 129–140.
- Stephen Roller, Emily Dinan, Naman Goyal, Da Ju, Mary Williamson, Yinhan Liu, Jing Xu, Myle Ott, Kurt Shuster, Eric Michael Smith, Y-Lan Boureau, and Jason Weston. 2020. Recipes for building an open-domain chatbot. *CoRR*.
- Alexis Ross, Ana Marasovic, and Matthew E. Peters. 2020. Explaining NLP models via minimal contrastive editing (mice). *CoRR*.
- Ashish Sharma, Inna W Lin, Adam S Miner, David C Atkins, and Tim Althoff. 2021. Towards facilitating empathic conversations in online mental health support: A reinforcement learning approach. In *WWW*.
- Tianxiao Shen, Tao Lei, Regina Barzilay, and Tommi S. Jaakkola. 2017. Style transfer from non-parallel text by cross-alignment. In *NIPS 2017, 4-9 December 2017, Long Beach, CA, USA*, pages 6833–6844.
- Kurt Shuster, Samuel Humeau, Antoine Bordes, and Jason Weston. 2020. Image-chat: Engaging grounded conversations. In *Proceedings of ACL 2020, Online, July 5-10, 2020*, pages 2414–2429.
- Karen Simonyan, Andrea Vedaldi, and Andrew Zisserman. 2014. Deep inside convolutional networks: Visualising image classification models and saliency maps. In *ICLR 2014, Banff, AB, Canada, April 14-16, 2014, Workshop Track Proceedings*.
- 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 ACL 2020, Online, July 5-10, 2020*, pages 2021–2030.
- Haoyu Song, Yan Wang, Weinan Zhang, Xiaojiang Liu, and Ting Liu. 2020. Generate, delete and rewrite: A three-stage framework for improving persona consistency of dialogue generation. In *Proceedings of ACL 2020, Online, July 5-10, 2020*, pages 5821–5831.
- Haoyu Song, Weinan Zhang, Yiming Cui, Dong Wang, and Ting Liu. 2019. Exploiting persona information for diverse generation of conversational responses. In *Proceedings of IJCAI 2019, Macao, China, August 10-16, 2019*, pages 5190–5196.

- Akhilesh Sudhakar, Bhargav Upadhyay, and Arjun Maheswaran. 2019. "transforming" delete, retrieve, generate approach for controlled text style transfer. In *Proceedings of EMNLP-IJCNLP 2019, Hong Kong, China, November 3-7, 2019*, pages 3267–3277.
- Sean Welleck, Jason Weston, Arthur Szlam, and Kyunghyun Cho. 2019. Dialogue natural language inference. In *Proceedings of ACL 2019, Florence, Italy, July 28- August 2, 2019, Volume 1: Long Papers*, pages 3731–3741.
- Jason Weston, Emily Dinan, and Alexander H. Miller. 2018. Retrieve and refine: Improved sequence generation models for dialogue. In *Proceedings of SCAI@EMNLP 2018, Brussels, Belgium, October 31, 2018*, pages 87–92.
- Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Rémi Louf, Morgan Funtowicz, Joe Davison, Sam Shleifer, Patrick von Platen, Clara Ma, Yacine Jernite, Julien Plu, Canwen Xu, Teven Le Scao, Sylvain Gugger, Mariama Drame, Quentin Lhoest, and Alexander M. Rush. 2020. Transformers: State-of-the-art natural language processing. In *Proceedings of EMNLP 2020 - Demos, Online, November 16-20, 2020*, pages 38–45.
- Thomas Wolf, Victor Sanh, Julien Chaumond, and Clement Delangue. 2019. Transfertransfo: A transfer learning approach for neural network based conversational agents. *CoRR*.
- Xing Wu, Tao Zhang, Liangjun Zang, Jizhong Han, and Songlin Hu. 2019a. Mask and infill: Applying masked language model for sentiment transfer. In *Proceedings of IJCAI 2019, Macao, China, August 10-16, 2019*, pages 5271–5277.
- Yu Wu, Furu Wei, Shaohan Huang, Yunli Wang, Zhoujun Li, and Ming Zhou. 2019b. Response generation by context-aware prototype editing. In *AAAI 2019, Honolulu, Hawaii, USA, January 27 - February 1, 2019*, pages 7281–7288.
- Yingce Xia, Fei Tian, Lijun Wu, Jianxin Lin, Tao Qin, Nenghai Yu, and Tie-Yan Liu. 2017. Deliberation networks: Sequence generation beyond one-pass decoding. 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 1784–1794.
- 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 ACL 2018, Melbourne, Australia, July 15-20, 2018, Volume 1: Long Papers*, pages 2204–2213.
- Weinan Zhang, Qingfu Zhu, Yifa Wang, Yanyan Zhao, and Ting Liu. 2019. Neural personalized response generation as domain adaptation. *World Wide Web*, 22(4):1427–1446.
- Peixiang Zhong, Chen Zhang, Hao Wang, Yong Liu, and Chunyan Miao. 2020. Towards persona-based empathetic conversational models. In *Proceedings of EMNLP 2020, Online, November 16-20, 2020*, pages 6556–6566.
- Chulun Zhou, Liangyu Chen, Jiachen Liu, Xinyan Xiao, Jinsong Su, Sheng Guo, and Hua Wu. 2020. Exploring contextual word-level style relevance for unsupervised style transfer. In *Proceedings of ACL 2020, Online, July 5-10, 2020*, pages 7135–7144.
- Hao Zhou, Minlie Huang, Tianyang Zhang, Xiaoyan Zhu, and Bing Liu. 2018a. Emotional chatting machine: Emotional conversation generation with internal and external memory. In *Proceedings of AAAI 2018, New Orleans, Louisiana, USA, February 2-7, 2018*, pages 730–739.
- Hao Zhou, Tom Young, Minlie Huang, Haizhou Zhao, Jingfang Xu, and Xiaoyan Zhu. 2018b. Commonsense knowledge aware conversation generation with graph attention. In *Proceedings of IJCAI 2018, July 13-19, 2018, Stockholm, Sweden.*, pages 4623–4629.
- Jun-Yan Zhu, Taesung Park, Phillip Isola, and Alexei A. Efros. 2017. Unpaired image-to-image translation using cycle-consistent adversarial networks. In *ICCV 2017, Venice, Italy, October 22-29, 2017*, pages 2242–2251.

A Annotation Guideline (Simplified)

The following guideline is provided to AMT crowdworkers when collecting reference responses:

“We aim at building human-like agents that have their own personal background. We need your help to correct some responses that are irrelevant to or contradictory to the speaker’s personal background. In each sample, you first see the dialogue history between the two speakers (Speaker1 and Speaker2), the original response by Speaker2, and the personal background of Speaker2. These background sentences are probably irrelevant to or contradictory to Speaker2’s response. Your task is to minimally edit the response such that it shows the background. Two requirements should be satisfied: 1) The edited response should show the personal background. 2) By “minimally edit” we mean that the edited response should maintain the contents in the response that are not contradictory to the background sentences. We have pasted the original response into the answer blank, and please edit it directly.”

B Human Evaluation Guidelines

We provide our human evaluation guidelines in `software.zip`, and we will make them pub-

lic. We briefly summarize the guidelines here, and more details can be found in `software.zip`.

B.1 Grounded Minimal Editing

To make our task more comprehensible to human participants, we reformulate our task as a *response correction* task. We define two types of mistakes made by a response: 1) *contradict* and 2) *ignore*. We first ask the participants **identify the type of mistake**, which encourages the participants to reason over persona-grounded dialogues. We specify two requirements to be satisfied by a *good* correction: 1) the mistakes are corrected, and 2) minimal changes, i.e., all words that are not contradictory to the expected personal background should be maintained, and if more than four of them are not maintained, then this requirement is not satisfied. Given two corrections *A* and *B*, participants **vote prefer A, none, or prefer B**. We ask them to choose *none* if neither *A* nor *B* is a good correction.

B.2 Transferability

For each response, we instruct the participants to answer four questions:

- **Knowledge (0 or 1)**. Does this response include world knowledge? 0: no; 1: yes. World knowledge includes *facts* and *commonsense* (see `software.zip` for details).
- **Empathy (0 or 1)**. Do you think this response is showing empathy? 0: no; 1: yes. Showing empathy means being aware of or being sensitive to the *feelings* or *experience* of the person being talked to (see `software.zip` for details).
- **Background occurrences**. There are two personal background sentences. How many of them are reflected by this response (examples omitted here)? Since we only care about if at least one of the personas are shown, **persona consistency** is **0** if the answer to this question is 0, and **1** if the answer is 1 or 2.
- **Grammaticality (0 or 1)**. Is this response grammatical? 0: no; 1: yes.

C Experimental Details

Models are evaluated on the validation set for every 500 steps, based on the Average metric. The batch size is 32. We use Adam (Kingma and Ba, 2015) with the initial learning rate 5×10^{-5} and gradient

clip 1.0. The learning rate decays by half when the Average metric does not improve for two validations, and training terminates after three decays. We detokenize the BPE tokens into English words for evaluation. More details for the Reproducibility Checklist are in `software.zip`.

D Model and Baseline Details

DeLorean-FT and our GME model use the GPT-2 (Radford et al., 2018) as the backbone, initialized by Huggingface Transformers (Wolf et al., 2020) checkpoint `gpt2` (DialogPT-small for DialogPT). UNMT and CycleGAN are Transformers with DistilGPT-2 encoder and decoder, initialized with `distilgpt2`. The auxiliary response-to-persona module in CycleGAN is implemented as a two-layer Transformer, initialized by the first two layers of DistilGPT-2.

E Data Samples and System Outputs

We provide several failure cases of GME in Table 7. Table 8-11 present some data samples and system outputs.

Dialogue history	Speaker 1: hey , i love vegas , playing the slots ! Speaker 2: vegas is fun . i do not eat much so i can stay skinny and wear those jeans . Speaker 1: yeah , i like to grill outside in the summer ! Speaker 2: over my skinny jeans i like to wear leggings . summer grilling is the best . Speaker 1: i also like mowing my lawn when the suns out .
Original response	Speaker 2: i put on the music when i mow the lawn or play video games .
Editing persona	i do not like music when i am working .
Human reference 1	me too , but i do not like music when i mow the lawn or play video games .
Human reference 2	i do not like to put on the music when i mow the lawn or play video games .
Human reference 3	i like to mow the lawn or play video games when suns out , but i do not like music when i am working .
GME (ours) // error	i put on the music when i mow the lawn or play music . i do not like music . // <i>Contradiction</i>
Context	Speaker 1: hello . how are you today ? Speaker 2: good today . just cooking some mexican food . i cooking but am not very good . Speaker 1: oh i like mexican food , but my favorite food are cheeseburgers . Speaker 2: i inherited some money and bought a yacht to travel , i try different foods traveling .
Original response	Speaker 1: i help out at a soup kitchen since i grew up poor .
Editing persona	i grew up with lots of cash .
Human reference 1	i grew up rich but i help out at a soup kitchen .
Human reference 2	i have lots of cash , but i help out at a soup kitchen .
Human reference 3	i help out at a soup kitchen since i grew up with lots of cash .
GME (ours) // error	i help out at a lot of cash since i grew up . i grew up with lots of cash . // <i>Grammatical error</i>
Dialogue history	Speaker 1: hello , how are you today ? Speaker 2: hello , i am fine thanks and you ? Speaker 1: i am good because i love music and play it all the time . Speaker 2: ah that is nice ! i play softball in my free time . Speaker 1: nice , trabajo is my favorite spanish word . Speaker 2: i wish i had time to learn another language , but i am busy with work . Speaker 1: yeah i want to study french next . Speaker 2: since i have been fired from my last job i have been working in insurance . Speaker 1: that is pretty cool ! i love to study spanish .
Original response	Speaker 2: i am a member of the army , served for 10 years now .
Editing persona	i am a school teacher , i teach middle school .
Human reference 1	i am a school teacher and teach middle school , served for 10 years now .
Human reference 2	i am a school teacher for 10 years now and i teach middle school .
Human reference 3	i am a middle school teacher and i have taught for 10 years .
GME (ours) // error	i am a teacher of the middle school . i teach middle school . // <i>Repetition</i>

Table 7: Failure cases

Dialogue history	<p>Speaker 1: hello , how are you doing ?</p> <p>Speaker 2: great , how are you ? i just finished watching one of my favorite documentaries . do you enjoy those ?</p> <p>Speaker 1: i am doing great , just tired . i just am unpacking boxes . i do not watch tv often .</p> <p>Speaker 2: did you just move ? i live here in pennsylvania with my husband .</p> <p>Speaker 1: yes , i bought my first house . i love pennsylvania , a lot of hills and very green .</p> <p>Speaker 2: good for you and congratulations on your new home !</p> <p>Speaker 1: thank you ! so what do you do for work ?</p> <p>Speaker 2: i just started working as a personal assistant about three months ago . how about you ?</p>
Original response	Speaker 1: that sounds fun , i am a teacher at the public school .
Editing persona	i work at a place that cleans cars .
Human reference 1	that sounds fun , i work at a place that cleans cars beside the public school .
Human reference 2	that sounds fun , i work at a place that cleans cars .
Human reference 3	that sounds fun , i am a teacher at the public school but i work as a car cleaner in part time .
UNMT (0.1)	that sounds fun , i am a teacher at the public school .
CycleGAN	that sounds fun , i am a teacher at the public school .
DeLorean-FT	
- $\gamma_{mix} = 0.75$	i work at a place that cleans cars .
- $\gamma_{mix} = 0.80$	i work at a place that cleans cars .
- $\gamma_{mix} = 0.85$	i work at a place that cleans cars .
GME (ours)	that sounds fun , i am a car mechanic at the place . i work at a place .

Table 8: Data sample and system outputs (*correction*)

Dialogue history	<p>Speaker 1: hi . i do not like working as a car salesman .</p> <p>Speaker 2: i recently broke my arm so i am not working .</p> <p>Speaker 1: what happened ? it is hard to do anything with a broken arm .</p> <p>Speaker 2: i blame my skateboarding friends .</p>
Original response	Speaker 1: do you think 40 is too old to go back to school ?
Editing persona	i am seventy two years old .
Human reference 1	do you think seventy two years old is too old to go back to school ?
Human reference 2	do you think seventy two is too old to go back to school ?
Human reference 3	do you think seventy years old is too old to go back to school ?
UNMT (0.1)	do you think 40 is too old to go back to school ?
CycleGAN	do you think 40 is too old to go back to school ? i am seventy seventy twelve years old .
DeLorean-FT	
- $\gamma_{mix} = 0.75$	do you skateboard ? i am a seventy two year old .
- $\gamma_{mix} = 0.80$	i am a seventy two year old man .
- $\gamma_{mix} = 0.85$	i am a seventy two year old man .
GME (ours)	do you think i am too old to go back to school ? i am seventy two years old .

Table 9: Data sample and system outputs (*correction*)

Dialogue history	Speaker 1: hello hope your having a good night tonight ? i know i am . Speaker 2: that is good ! i am watching football highlights . Speaker 1: really what teams . and who are you rooting for . Speaker 2: i am a green bay packer fan . do you like football ? Speaker 1: no i am not a fan but my son is .
Original response	Speaker 2: how many kids do you have ?
Editing persona	i don t like kids .
Human reference 1	i don t like kids very much . how many kids do you have ?
Human reference 2	i don t have kids . how many kids do you have ?
Human reference 3	how many kids do you have ? i don t like kids .
UNMT (0.1)	how many kids do you have ?
CycleGAN	how many kids do you have ?
DeLorean-FT	
- $\gamma_{mix} = 0.75$	how many kids do you have ?
- $\gamma_{mix} = 0.80$	i am a kid . i do not like kids .
- $\gamma_{mix} = 0.85$	i am a kid . i do not like kids .
GME (ours)	how many kids do you have ? i do not like kids .

Table 10: Data sample and system outputs (*add new information*)

Dialogue history	Speaker 1: do you have any kids ? i have one in middle school . Speaker 2: i do not . my husband and i have chosen to not have any .
Original response	Speaker 1: probably a smart decision , too many people on the planet .
Editing persona	i recently started to work online .
Human reference 1	probably a smart decision , i recently started to work online because too many people on the planet .
Human reference 2	probably a smart decision , too many people on the planet that is why i recently started to work online .
Human reference 3	probably a smart decision , too many people on the planet . i recently started to work online .
UNMT (0.1)	probably a smart decision , too many people on the planet .
CycleGAN	probably a smart decision , too many people on the planet .
DeLorean-FT	
- $\gamma_{mix} = 0.75$	i am a computer science major . i am currently working online .
- $\gamma_{mix} = 0.80$	i am a computer science major . i am currently working online .
- $\gamma_{mix} = 0.85$	i am a computer science . i am currently working online .
GME (ours)	probably a smart decision , too many people on the planet . i am working online now .

Table 11: Data sample and system outputs (*add new information*)