Persona Expansion with Commonsense Knowledge for Diverse and Consistent Response Generation

Donghyun Kim[1] **Youbin Ahn**[1] **Wongyu Kim**[1] **Chanhee Lee**[1] **Kyungchan Lee**[1]
Kyong-Ho Lee[1] **Jeonguk Kim**[2] **Donghoon Shin**[2] **Yeonsoo Lee**[2]
[1]Department of Computer Science, Yonsei University
[2]NCSOFT
{dhkim92,ybahn,rladnjsrb9999,heederer,steven800,khlee89}@yonsei.ac.kr
{jukim,dhshin,yeonsoo}@ncsoft.com

Abstract

Generating diverse and consistent responses is the ultimate goal of a persona-based dialogue. Although many studies have been conducted, the generated responses tend to be generic and bland due to the personas' limited descriptiveness. Therefore, it is necessary to expand the given personas for more attractive responses. However, indiscriminate expansion of personas threaten the consistency of responses and therefore reduce the interlocutor's interest in conversation. To alleviate this issue, we propose a consistent persona expansion framework that improves not only the diversity but also the consistency of persona-based responses. To do so, we define consistency criteria to avoid possible contradictions among personas as follows: 1) Intra-Consistency and 2) Inter-Consistency. Then, we construct a silver profile dataset to deliver the ability to conform with the consistency criteria to the expansion model. Finally, we propose a persona expansion model with an encoder-decoder structure, which considers the relatedness and consistency among personas. Our experiments on the Persona-Chat dataset demonstrate the superiority of the proposed framework.

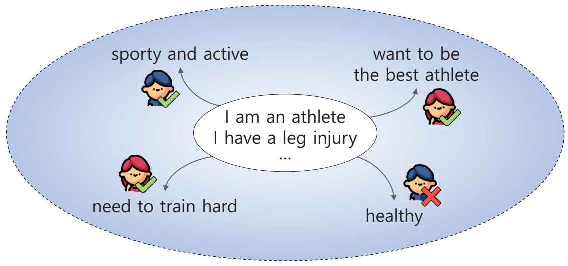

Figure 1: An example of persona expansion with commonsense. The center circle is the predefined profile, and the dashed circle is the valid expandable range.

1 Introduction

In the field of open-domain dialogues, maintaining long-term interaction with users by generating human-like responses has been a persistent goal. As efforts along this line of research, persona-based dialogue models aim at generating more engaging and consistent responses based on personal traits (Li et al., 2016b; Qian et al., 2018). As the most widely used benchmark dataset, Persona-Chat (Zhang et al., 2018) consists of personal profiles, which contain 3-5 persona sentences (hereinafter referred to as "personas"), and dialogues conditioned on personas. With the Persona-Chat, many researches have attempted to improve the quality of responses from various perspectives (Song et al., 2020; Liu et al., 2020; Song et al., 2021; Cao et al., 2022). However, since the descriptiveness of personas is limited, they have difficulties in converting what personas imply into attractive responses.

Unlike dialogue models, humans converse beyond the explicit persona itself through rational reasoning with commonsense knowledge. The commonsense reasoning ability is essential for making an interlocutor engage in conversations with a dialogue agent. Figure 1 illustrates an example of persona expansion through commonsense reasoning. Given the persona '*I am an athlete*', it is possible to infer the characteristic '*sporty and active*' with high probability. The implications of '*want to be the best athlete*' and '*need to train hard*' are also predictable. In order to mimic human reasoning ability, previous works expand original personas (Majumder et al., 2020) or infer new personas from dialogue utterances (Kim et al., 2022). However, the consistency among expanded personas was not explicitly considered. To prolong a conversation, both the engagingness and the consistency of responses are very important. Therefore, we define the two consistencies that must comply with while improving the diversity of responses as follows: 1) **Intra-Consistency** and 2) **Inter-Consistency**. First, the Intra-Consistency means that expanded personas must not contradict the original persona.

Proceedings of the 17th Conference of the European Chapter of the Association for Computational Linguistics, pages 1139–1149
May 2-6, 2023 ©2023 Association for Computational Linguistics

If the response conditioned on the expanded persona contradicts the original persona, the interlocutor loses interest in the conversation. Second, the Inter-Consistency means that expanded personas must not contradict other personas in the same profile. For instance, as shown in Figure 1, the attribute of '*healthy*' can be expanded from the original persona of '*I am an athlete*'. However, considering the other persona in the profile, '*I have a leg injury*', this comes to an inappropriate expansion. Therefore, it is necessary to follow the Intra- and Inter-Consistency when expanding personas to improve the persona-based responses.

To achieve the above goals, we propose a novel **P**ersona **E**xpansion framework with **C**ommonsense **K**nowledge (**PECK**). PECK aims to not only improve the diversity of responses but also avoid the problem of contradictions among personas. Both abilities to understand the relatedness and to reason with commonsense knowledge are required to satisfy the consistency criteria. Unfortunately, there are no supervised training datasets for persona expansion. Furthermore, it costs a lot to annotate them to train both abilities. Therefore, we build a silver profile dataset (see Section 4) similar to the gold profile in the Persona-Chat, with simple manipulation of the commonsense knowledge graph ANION (Jiang et al., 2021). In addition, we propose a persona expansion model (see Section 5) consisting of an encoder and a decoder for consistent expansion. The bidirectional encoder learns to understand the relatedness among personas within the same profile. Then, the auto-regressive decoder learns to expand personas consistently with negative log-likelihood (NLL) and unlikelihood (UL) (Welleck et al., 2019a) training objectives. Finally, we conduct extensive experiments on persona expansion and persona-based response generation. Experimental results confirm the superiority of the proposed consistent expansion framework.

The contributions of this work are as follows:

- A novel persona expansion framework that satisfies essential consistency criteria is proposed for attractive chit-chat responses.

- We generate consistent personas with a manipulated silver profile dataset and unlikelihood training objective.

- We show the effectiveness of the proposed expansion framework via experimental results.

2 Related Work

2.1 Persona Expansion with Commonsense

COMPAC proposed by Majumder et al. (2020) makes a fine-grained choice of a persona to generate more engaging chit-chat responses. They expand the predefined personas leveraging COMET (Bosselut et al., 2019) trained with the commonsense knowledge of ATOMIC (Sap et al., 2019). For instance, COMET deduces expanded personas such as '*I am caring*, *I want to adopt a cat*' from the original persona '*I love animals*' along the nine relations in ATOMIC. Meanwhile, Kim et al. (2022) propose the dual task framework of debiasing persona-based dialogues via a data-driven approach. Similary, they utilize COMET to perform commonsense-based expansion before matching personas with utterances in a dialogue. The commonsense-based expansion in the conventional methods improves the engagingness of persona-based dialogues. However, they did not consider that consistency should be kept in the expansion process. To address this issue, we propose a novel approach of expanding personas and keeping the proposed Intra- and Inter-Consistency among them.

2.2 Commonsense Knowledge Graphs

As the demand for the representation of relations between real-world events, some recent studies on event commonsense knowledge graphs (CSKGs) have been conducted. ATOMIC (Sap et al., 2019) is an event-based CSKG, containing textual descriptions of inferential knowledge on real-world events. ATOMIC provides nine if-then causal relations, in which the reasoning such as causes and effects focuses on the agent (i.e., xAttr, xEffect, xIntent, xNeed, xReact, xWant.) or the other (i.e., oEffect, oReact, oWant). As a follow-up study, Hwang et al. (2021) propose ATOMIC-2020, which contains social and physical relations as well as eventive ones. Further, Jiang et al. (2021) introduce ANION, which focuses on negated and contradictory events contrary to the affirmative ones of ATOMIC. ANION contains 624k if-then rules consisting of negated and contradictory relations. We construct a silver profile dataset by leveraging the inter-event inferential commonsense of ANION via simple manipulation. And then, we compose positive and negative samples of unlikelihood training (Welleck et al., 2019a) for consistent persona expansion by leveraging the oppositions between each profile in the dataset.

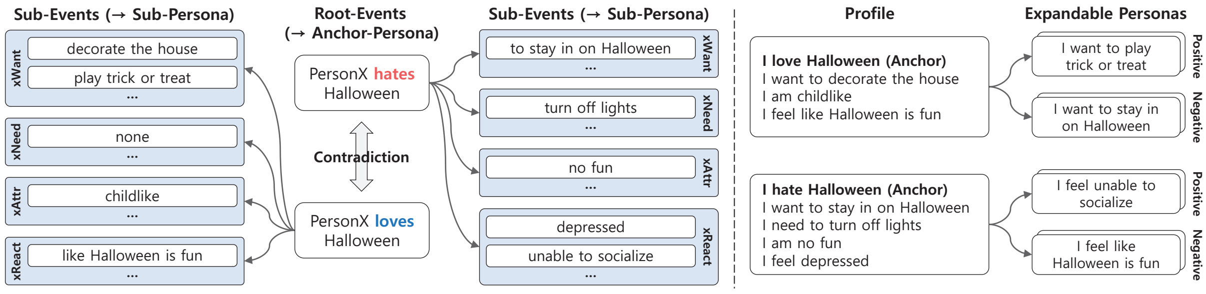

(1) Examples of contradictory root-events and their sub-events

(2) Positive and Negative Training Samples

Figure 2: (1) An example of contradictory root-events and their sub-events in ANION. The root-event and its corresponding sub-events are designated anchor-persona and sub-personas, respectively. (2) Training samples for consistent persona expansion. Reasonable expansions are classified as positive samples E^+, and contradictories are classified as negative samples E^-.

3 Problem Statement

In this paper, our goal is to generate expanded personas, which contribute to improving the quality of persona-based responses. To do so, the consistency criteria that expanded personas must comply with are as follows: 1) **Intra-Consistency** and 2) **Inter-Consistency**. Formally, given an original persona $p_i \in P$, where $P = \{p_1, p_2, ..., p_n\}$ is the predefined profile, a persona expansion model aims to generate expanded personas $E_i = \{e_1, e_2, ..., e_m\}$, which do not contradict the original persona p_i and all the remaining personas in profile P. For the model to learn these consistencies, we need original-expandable persona pairs with a positive or negative label, and the profile where the original persona belongs. That is to say, the training dataset should contain {*Original Personas* (*Profile*), *Expandable Personas*, *Label*}. Regrettably, there are no training datasets of this form, and human annotation is expensive. Hence, we construct a new dataset, silver profile, appropriate for our expansion goal through simple manipulation on ANION, the existing commonsense knowledge graph. The detailed construction process will be described in Section 4. Finally, our model trained with the silver profile aims to generate expanded personas consistently from gold profiles within Persona-Chat.

4 Data Construction

Leveraging ANION, we construct a new dataset for training consistent persona expansion. Through exploratory analysis on Persona-Chat, we identified some prominent features of profiles. First, each profile is mostly made up of personas, which are correlated or at least non-contradictory to one another. In the gold profile example in Table 1, three of the five personas (1, 2, 3) are directly correlated, and the others are slightly related. Second, most of all personas indicate the speaker's preferences (i.e., I like –, I love –, I want –) or states (i.e., I am –). Based on these structural and semantical features of profiles, we build a silver profile dataset similar to the gold profile in Persona-Chat to train our expansion model. Statistics of the silver profile dataset are given in Appendix B.

Gold Profile	Silver Profile
1. I play football. 2. My position is linebacker. 3. I am an athlete. 4. I had 128 tackles last year. 5. My team is the baltimore ravens.	1. I love Halloween. 2. I want to decorate the house. 3. I am childlike. 4. I feel like Halloween is fun.

Table 1: Examples of the gold profile from Persona-Chat and silver profile from ANION.

4.1 Profilizing Commonsense Events

With reference to the aforementioned feature of the profile, the *root-event* existing in ANION and the *sub-events* connected by if-then relations can be defined as the speaker's profile. In this case, the root-event selected for profilizing is the sentence described only for *PersonX*, not including *PersonY*.[1] As shown on the left of Figure 2, the root-event is defined as a dominant *anchor-persona* in the profile, and each sub-event is defined as a *sub-persona* that can be inferred from the anchor-persona. To filter mislinked relations out during ANION deployment, we leverage the RoBERTa (Liu et al., 2019) model fine-tuned with the MNLI

[1]In ANION, there are events related to the agent (i.e., *PersonX* –) and events related to the agent and the other people (i.e., *PersonX* – *PersonY* –). In this study, events with *PersonY* are excluded because we only handle personas representing the traits of agent '*I*'.

(Williams et al., 2018) dataset to verify the non-contradiction between a root-event and its sub-events. Also, based on the fact that most personas represent the speaker's preferences or states, we get sub-events from four if-then relations as follows: *xWant*, *xNeed*, *xAttr*, *xReact*. Then, we change the subject of each event to '*I*' and conduct subject-verb matching to convert them into a sentence of a persona type, such as <*PersonX loves Halloween*> $\rightarrow$ '*I love Halloween*'. An example of silver profile is shown on the right side of Table 1. If '*I love Halloween*' is the anchor-persona, then '*I want to decorate the house*' and '*I feel like Halloween is fun.*' are designated as sub-personas.

4.2 Assigning Expandable Personas to Profile

A silver profile has a maximum of five personas; one is the anchor-persona, and the others are sub-personas from four if-then relations. Since there are one or more sub-personas for each relation, sentences not designated as sub-personas can be assigned as expandable personas. As shown on the right side of Figure 2, '*I want to play trick or treat*' is not designated as a sub-persona, so it is treated as an expandable persona. Through this process, the expandable personas that can be inferred for the four relations are assigned to the profile.

4.3 Aligning Contradictory Profiles

Inspired by Li et al. (2020), we leverage NLL and UL objectives which require both positive and negative samples to endow the model with the consistent expansion ability. Therefore, we align the corresponding contradictory profiles to construct positive and negative samples. In ANION, negated root-events such as <*PersonX hates Halloween*> and <*PersonX does not loves Halloween*> contradict their affirmative <*PersonX loves Halloween*>. Using these opposites as pivots, we construct negative samples by swapping the expandable personas of each other. As shown on the right side of Figure 2, for profiles with '*I love Halloween*' as an anchor-persona, '*I want to play trick or treat*' is a positive expandable persona, and '*I want to stay in on Halloween*' comes to a negative one.

Definition 1 *A silver profile dataset is defined as* $D_{sp} = \{(P_i, E_i^{+/-})\}_{i=1}^{K}$, *where* P_i *is a synthetic profile,* $E_i^+ = \{E_i^{want}, E_i^{need}, E_i^{attr}, E_i^{react}\}$ *is a set of expandable personas for four relations,* E_i^- *is a set of contradictories, and each* $E_i^{rel} = \{e_1^{rel}, e_2^{rel}, ..., e_{m_{rel}}^{rel}\}$ *consists of* m_{rel} *personas.*

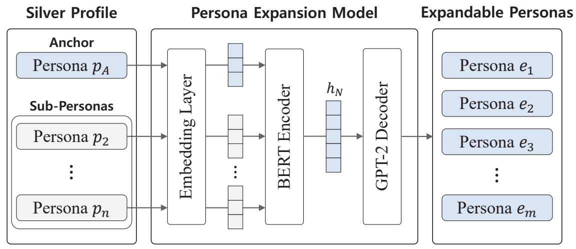

Figure 3: Overview of the proposed expansion model. In the training phase, the model leverages the silver profile dataset to learn consistencies among personas.

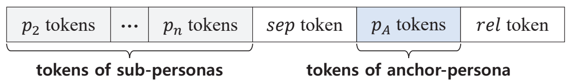

Figure 4: Input token template for training the persona expansion model.

5 Consistent Persona Expansion

5.1 Overview

In this section, we propose a novel consistent persona expansion model. To achieve the consistency criteria mentioned above, we design the model with an encoder-decoder structure as illustrated in Figure 3. The encoder learns to understand the relatedness among personas in the profile, and the decoder learns to consistently reason expandable personas with commonsense knowledge. This separated architecture allows the model to generate expanded personas satisfying the consistencies.

Formally, given a silver profile dataset D_{sp}, the model takes the profile $P = \{p_1, p_2, ..., p_n\}$ as input, where p_1 $(= p_A)$ is the anchor-persona, and the others are sub-personas. Then, the target expandable personas are $E^{rel} = \{e_1^{rel}, e_2^{rel}, ..., e_{m_{rel}}^{rel}\}$.

5.2 Encoder

The encoder is initialized by the pre-trained BERT (Devlin et al., 2019) embedding to inherit its bidirectional language understanding ability.

The input consists of unfolded persona profile P to recognize a correlation among personas in the profile. As illustrated in Figure 4, a separate token $[sep]$ is placed between the sequence of sub-personas and the anchor-persona, and a relation token $[rel]$ is placed at the end of the input. The formatted input is as follows:

$$p_2, ..., p_n, [sep], p_A, [rel]. \tag{1}$$

Then, the embedding layer transforms each token of the input into a vector representation. At this time, the sum of the corresponding token and position embedding is used as the vector representation. The resulting vector representation is denoted as $input$. Next, the encoder takes the vector representation and performs multi-head attention (Vaswani et al., 2017) as follows:

$$h_{i+1} = FFN(MultiHead(h_i, h_i, h_i)), \quad (2)$$

where $MultiHead(query, key, value)$ calculates the scaled dot-product attention of query, key, value, $FFN(\cdot)$ is a fully connected feed-forward network, and $h_0 = input$. The encoder outputs the vector representation h_L after L identical layers. The vector representation h_L implies context information for expansion and the correlatedness among all personas in the profile.

5.3 Decoder

The decoder is initialized by the pre-trained GPT-2 (Radford et al., 2019) to leverage its robust auto-regressive generation ability.

Unlike the encoder, the decoder adopts the cross-attention and left-to-right mask. First, the cross-attention is performed to reflect the context information among personas. In the cross-attention, the query is the vector representation from previous layer, and the output vector h_L of the encoder is given as the key and value:

$$v_{i+1} = FFN(MultiHead(v_i, h_L, h_L)), \quad (3)$$

where v_0 is initialized from the embeddings of the decoder input $e_j^{rel} (\in E^{rel})$ for given relation rel in Equation 1. Then, the decoder outputs the vector representation v_L after L identical layers. Finally, after the linear and softmax layer, we get the probabilities of each token in the output sequence.

5.4 Training Objectives

For the consistent persona expansion, we exploit negative log-likelihood loss (NLL) and unlikelihood loss (UL). NLL is a conventional objective function for generation tasks that maximizes the probability of the next target token. In contrast, UL is an objective function that minimizes the probability of the next target token, usually used to prevent the model from generating dull and repetitive sequences. In our setting, referred to Li et al. (2020), we use NLL and UL for positive samples (D_{sp}^+) and negative samples (D_{sp}^-), respectively. The expression for each sample is as follows:

$$D_{sp}^+ = \{P, E^+\}, D_{sp}^- = \{P, E^-\}, \quad (4)$$

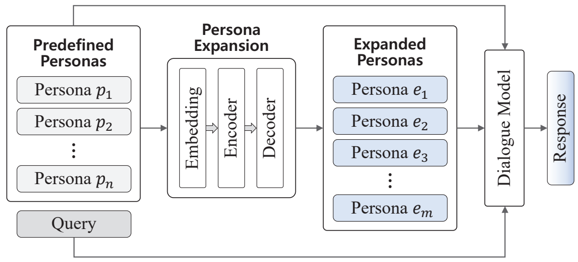

Figure 5: Response generation with persona expansion.

where P is the profile, E^+ is the expandable personas, and E^- is the contradictory personas. For simplicity of expression, we omit the index notation i. See Definition 1 for the exact expression. The NLL and UL losses are calculated as follows:

$$L_{NLL}^{D_{sp}^+} = -\sum_{i=1}^{|e|} log(p_\theta(t_i|P, e_{<i})), \quad (5)$$

$$L_{UL}^{D_{sp}^-} = -\sum_{i=1}^{|e|} log(1 - p_\theta(t_i|P, e_{<i})), \quad (6)$$

where e is each persona in E^+/E^-, and t_i is the token of the persona e at time step i. The final objective function of the model computes the weighted sum of NLL and UL losses as follows:

$$Loss = \alpha L_{NLL}^{D_{sp}^+} + \beta L_{UL}^{D_{sp}^-}, \quad (7)$$

where α and β are hyperparameters. The optimized objective function is obtained with iterative loss back-propagation.

5.5 Persona Expansion

Through the training with silver profile, our model acquires the ability to understand the relatedness among input personas and to generate consistent expanded personas. The model is leveraged to expand the personas of the gold profile in the Persona-Chat. For each persona in the gold profile, we generate five expanded personas per relation (i.e., xWant, xNeed, xAttr, xReact) by beam-search decoding. Then, the number of expanded personas per an original persona comes to 20 (= 5 × 4).

5.6 Response Generation

After persona expansion, we generate diverse and consistent responses utilizing the original and expanded personas, as illustrated in Figure 5. In this paper, we leverage two dialogue models to prove the effectiveness of the proposed framework on response generation.

GPT-2 (Wolf et al., 2019): A GPT-2 model simply concatenates all personas with dialogue history. In our setting, up to five expanded personas with the highest NLI scores on the predefined profile are selected and fed into the dialogue model.

COMPAC (Majumder et al., 2020): A dialogue model that generates fine-grained persona-based responses through expansion and selection. In our setting, the expanded personas generated by PECK framework are selected by COMPAC and used for response generation.

6 Experiments

6.1 Dataset

We conducted experiments on both persona expansion and response generation with Persona-Chat, converted for the ConvAI2 benchmark (Dinan et al., 2020). The Persona-Chat consists of dialogues between two speakers, each utterance conditioned on personas in the speaker's profile. We generated the expanded personas using personas in the predefined profile. And we carried out the response generation using utterances and expanded personas. In our experiments, 16,878 instances were used for training, 1,000 for validation and 1,000 for testing.

6.2 Baselines

To verify the effectiveness of the constructed silver profile dataset and the proposed expansion model, the following three baseline models were adopted and compared.

1. **COMET$_{\textbf{ATOMIC}}$** (Bosselut et al., 2019): As a GPT-2 model fine-tuned on ATOMIC, it expands the original personas along the four relations mentioned above for comparison.

2. **COMET$_{\textbf{SP w/o UL}}$**: We fine-tuned the GPT-2 model with a silver profile dataset as a comparative model, which takes a concatenated sequence of personas as input. UL objective was not used in training.

3. **COMET$_{\textbf{SP w/ UL}}$**: We fine-tuned the GPT-2 using a silver profile dataset and UL with negative samples. The other settings except UL are the same as the above model.

6.3 Evaluation Metrics

6.3.1 Automatic Evaluation

Four metrics were utilized for automatic evaluation. We employed Perplexity (**PPL**), which measures how well a model predicts response tokens, where lower means better. We employed Distinct-1/2 (**Dist-1/2**) (Li et al., 2016a), which measures the diversity of sentences via calculating the ratio of distinct words against total uni/bi-grams. We employed F1 score of BERTScore (**F-BERT**) (Zhang et al., 2019), which measures how semantically relevant the expanded persona is to the original persona or generated response is to the selected persona. We employed Consistency Score (**C.Score**) (Madotto et al., 2019), which measures the consistency of the expanded persona (or response) for the predefined profile P. The scoring process is as follows:

$$NLI(p_i, e) = \begin{cases} 0, & \text{if } e \text{ contradicts } p_i, \\ 1, & \text{otherwise,} \end{cases} \tag{8}$$

$$C.Score(e) = \sum_{i=1}^{n} NLI(p_i, e), \tag{9}$$

where p_i is each persona in the profile P, and e is the expanded persona (or response) to score. RoBERTa was employed for the NLI model, which is fine-tuned with the DNLI (Welleck et al., 2019b) dataset. The Consistency Score is calculated, and scaled to a value between 0 and 1.

6.3.2 Human Evaluation

For qualitative assessment, we conducted evaluations by three annotators. We recruited annotators proficient in the language from an external agency. Annotators had sufficient knowledge of the domain but knew nothing of the proposed model. During the evaluation, there was no interaction among annotators that could affect each other's ratings. Also, annotators were properly compensated for the work they performed.

We randomly selected 100 samples each to evaluate persona expansion and response generation. Human annotators were asked to evaluate samples with the following four metrics: (1) Engagement, (2) Relevance, (3) Diversity, (4) Intra-Consistency, and (5) Inter-Consistency. Ratings range from 1 to 5, with higher scores indicating better. Intra-Consistency measures the consistency between the expanded persona and the original persona. Inter-Consistency is scaled to a value between 1 to 5, after summing up the one-to-one Intra-Consistency score between the expanded persona and each persona in the profile. Further, we conducted human A/B test that provide consensus among annotators for consistency.

Model	Automatic Evaluation				Human Evaluation			
	Dist-1	Dist-2	F-BERT	C.Score	Engagement	Relevance	Intra-Con.	Inter-Con.
$\text{COMET}_{\text{ATOMIC}}$	0.041	0.162	0.806	0.54	2.92	3.01	3.04	2.24
$\text{COMET}_{\text{SP w/o UL}}$	0.044	0.165	0.810	0.66	3.08	3.06	3.12	2.67
$\text{COMET}_{\text{SP w/ UL}}$	0.043	0.164	0.818	0.70	3.05	3.18	3.58	3.52
$\text{PECK}_{\text{w/o UL}}$	**0.052**	**0.183**	<u>0.822</u>	<u>0.74</u>	**3.12**	<u>3.25</u>	<u>3.61</u>	<u>3.73</u>
$\text{PECK}_{\text{w/ UL}}$	<u>0.047</u>	<u>0.179</u>	**0.835**	**0.83**	<u>3.10</u>	**3.70**	**4.11**	**4.05**

Table 2: Automatic and human evaluation results on the persona expansion. The best results are **bolded**, and the second-best are <u>underlined</u>.

6.4 Persona Expansion Results

The automatic and human evaluation results are reported in Table 2. Compared with baselines, PECK outperformed other expansion models in all metrics regarding diversity and consistency. In particular, relevance (i.e., F-BERT, Relevance) and consistency (i.e., C.Score, Intra/Inter-Consistency) were significantly improved compared to diversity (i.e., Dist-1/2, Engagement). This result demonstrated the effectiveness of the proposed framework for consistent persona expansion.

6.4.1 Anlalysis

Through a detailed analysis of the results, we identified the contributions of the following three implements: 1) silver profile dataset, 2) encoder-decoder structure, and 3) UL objective.

First, from the results of $\text{COMET}_{\text{ATOMIC}}$ and $\text{COMET}_{\text{SP w/o UL}}$ in Table 2, the latter model using the silver profile dataset for training showed better performance in all evaluation metrics. Especially, C.Score showed a higher improvement compared to other automatic metrics. In human evaluation, Inter-Consistency increased more significantly than other metrics. This result showed that the presence of the silver profile helped improve the consistency of the expanded personas. In other words, it means that the silver profile sufficiently imbued the expansion model with the characteristics of the profile and correlatedness among personas.

To analyze the impact of the encoder-decoder structure, COMET_{SP} and PECK were trained using the same silver profile dataset. Different from PECK, COMET_{SP} is a model based on GPT-2 with a decoder-only structure. Our experimental results showed that PECK outperformed COMET_{SP} despite being trained on the same dataset. In addition, we found significant increases in C.Score and Inter-Consistency, the metrics representing one-to-many consistency. This means that the encoder has effectively fused personas within the profile so that the decoder could generate expanded personas consistently for the given profile.

PECK vs. $\text{COMET}_{\text{ATOMIC}}$	Intra-Consistency			Inter-Consistency		
	Win	Lose	k	Win	Lose	k
xWant	72.4	11.6	0.62	76.1	9.4	0.71
xNeed	68.3	18.4	0.52	72.7	10.2	0.66
xAttr	74.5	10.8	0.64	79.5	7.8	0.73
xReact	61.8	23.7	0.48	69.4	14.6	0.54

Table 3: Human A/B results for consistency. Ties are not indicated in the table. The values of Fleiss' kappa k (Fleiss, 1971) for all results are in $0.4 < k < 0.8$, indicating moderate and substantial agreements.

In terms of UL objective, both $\text{COMET}_{\text{SP w/ UL}}$ and $\text{PECK}_{\text{w/ UL}}$ models showed higher relevance and consistency than those without UL training. We found that Dist-1/2 and Engagement metrics slightly decreased when using UL training. This result seems to be because the constraint of generation by UL training defected the range of expanded personas. However, despite the decrease in diversity, it can be judged to be a sufficiently attractive expansion. The analysis demonstrated how effective the UL training is in improving consistency while keeping diversity.

6.4.2 Human A/B Test

To compare the consistency of the expansion for each relation, we conducted a human A/B test. Human annotators were asked to choose a more consistent persona between personas expanded by $\text{COMET}_{\text{ATOMIC}}$ and PECK. Table 3 shows that Intra- and Inter-Consistency of the persona expanded by PECK outperformed in all relations. In particular, the performances for Inter-Consistency were better than Intra-Consistency. Through these results, we found that PECK is effective in solving Inter-Consistency contradiction that the existing model cannot control.

Dialogue Model	Expansion Model	Automatic Evaluation				Human Evaluation			
		PPL	Dist-1	Dist-2	C.Score	Engagement	Diversity	Relevance	Consistency
GPT-2	-	22.12	0.054	0.108	0.45	2.08	2.64	2.98	2.25
	$\text{COMET}_{\text{ATOMIC}}$	21.43	0.206	0.297	0.39	3.18	2.78	2.71	2.12
	$\text{PECK}_{\text{w/ UL}}$	21.18	0.223	0.345	<u>0.67</u>	3.22	3.31	3.05	<u>3.24</u>
COMPAC	-	19.84	0.146	0.253	0.58	2.31	3.02	<u>3.14</u>	2.86
	$\text{COMET}_{\text{ATOMIC}}$	<u>16.36</u>	<u>0.836</u>	<u>0.964</u>	0.64	<u>3.35</u>	<u>3.42</u>	3.07	2.97
	$\text{PECK}_{\text{w/ UL}}$	**16.21**	**0.851**	**0.986**	**0.76**	**3.43**	**3.67**	**3.22**	**3.53**

Table 4: Automatic and human evaluation results on the response generation. The best results are **bolded**, and the second-best are <u>underlined</u>.

6.5 Response Generation Results

The automatic and human evaluation results on the response generation are reported in Table 4. The dialogue models with expansion showed better performance than those without expansion. Especially, COMPAC with PECK outperformed the others in all evaluation metrics. In the following sections, we further analyze the diversity and consistency of generated responses.

6.5.1 Analysis on Diversity

The results of the response generation showed that persona expansion greatly contributes to improving the diversity of responses. In automatic evaluation, Dist-1/2 metrics significantly outperformed when using expansion models than when not using them. In particular, the values of Dist-1/2 showed greater increases when using the personas expanded by PECK. The impact of expanded personas was noticeable in human evaluation. In Engagement and Diversity, the responses based on expanded personas were rated well by human annotators. This means that PECK framework is helpful in generating attractive responses, which are essential to engage interlocutors and prolong a conversation.

6.5.2 Analysis on Consistency

In terms of consistency, the performance increased significantly when PECK framework was applied. Further, PECK outperformed COMET on both dialogue models. When COMET expansion was utilized for response generation, C.Score, Relevance, and Consistency metrics were, in some cases, lower than without expansion. GPT with PECK achieved better C.Score and Consistency than COMPAC with COMET, despite not performing fine-grained persona selection. This suggests that COMET persona expansion help improve diversity but adversely affects consistency. On the other hand, PECK not only improved the diversity of responses but was also the best in consistency. These results demonstrated that the proposed expansion framework is effective at generating diverse and consistent responses.

Intra-Consistency	
Original Persona	I am primarily a **meat eater**
$\text{COMET}_{\text{ATOMIC}}$	I want to **avoid meat** I am **vegetarian** I feel satisfied
$\text{PECK}_{\text{w/ UL}}$	I want to **eat meat** I am **carnivore** I feel happy
Inter-Consistency	
Original Profile	I love to rest at home on the weekend (A) I am a **hard worker** My hobby is walking with my dog
$\text{COMET}_{\text{ATOMIC}}$	I am **lazy** I want to take a rest
$\text{PECK}_{\text{w/ UL}}$	I am **tired** I want to relax

Table 5: Examples of expanded personas. Each sample was picked from all expansions. Red indicates expansion that contradicts the original.

6.6 Case Study on Intra-/Inter-Consistency

Table 5 shows sample cases of persona expansion on Intra-/Inter-Consistency. In Intra-Consistency case, given the original persona '*I am primarily a meat eater*', COMET generated contradictory personas like '*I want to avoid meat*' and '*I am vegetarian*'. On the other hand, PECK generated consistent personas. In Inter-Consistency case, '*I love to rest at home on the weekend*' is an anchor-persona, and '*I am a hard worker*' is a sub-persona to be considered when expanding. COMET generated expanded personas such as '*I am lazy*' and '*I want to take a rest*'. Here, the expanded persona '*I am lazy*' contradicts the contents of the sub-persona. On the contrary, the personas expanded by PECK are consistent with the original profile.

7 Conclusions and Future Work

In this paper, we proposed a novel persona expansion framework to improve the diversity and consistency of the persona-based dialogue. Leveraging the existing commonsense knowledge graphs, we built a silver profile dataset to deliver the ability to conform with the consistency criteria (Intra-/Inter-Consistency) to the expansion model. The expansion model of the encoder-decoder structure was trained using the silver profile dataset with unlikelihood training. After training, the model generated expanded personas that satisfy the consistency criteria. The experimental results proved that the proposed persona expansion is effective at improving the diversity of responses while keeping their consistency with personas. In the future, we plan to extend PECK to leverage more external knowledge, such as ConceptNet (Speer et al., 2017).

Limitations

In this study, the expansion target was limited to personas with a first-person subject. However, if other forms of persona expansion are considered, more diverse and richer responses can be generated. To solve this problem, we plan to research strategies to more effectively utilize various existing commonsense datasets in the persona-based dialogue model.

Ethical Considerations

Through our analysis of the persona expansion, we found that expanded personas highly depend on the commonsense knowledge used in the model training. We built a silver profile dataset utilizing the commonsense knowledge graph ANION and our expansion model trained with the silver profile dataset. As mentioned in ANION, linguistic, social, and cultural biases may exist within the ANION, the source of our silver profile dataset. These biases or stereotypes will affect the meaning of the personas expanded by our framework. This study did not discuss this aspect, but we believe it is an issue worth considering in future works.

Acknowledgements

We thank all anonymous reviewers for their valuable comments on this work. This work was supported by NCSOFT NLP Center. Kyong-Ho Lee is the corresponding author.

References

Antoine Bosselut, Hannah Rashkin, Maarten Sap, Chaitanya Malaviya, Asli Celikyilmaz, and Yejin Choi. 2019. COMET: Commonsense transformers for automatic knowledge graph construction. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 4762–4779, Florence, Italy. Association for Computational Linguistics.

Yu Cao, Wei Bi, Meng Fang, Shuming Shi, and Dacheng Tao. 2022. A model-agnostic data manipulation method for persona-based dialogue generation. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 7984–8002, Dublin, Ireland. Association for Computational Linguistics.

Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. BERT: Pre-training of deep bidirectional transformers for language understanding. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4171–4186, Minneapolis, Minnesota. Association for Computational Linguistics.

Emily Dinan, Varvara Logacheva, Valentin Malykh, Alexander Miller, Kurt Shuster, Jack Urbanek, Douwe Kiela, Arthur Szlam, Iulian Serban, Ryan Lowe, et al. 2020. The second conversational intelligence challenge (convai2). In *The NeurIPS'18 Competition: From Machine Learning to Intelligent Conversations*, pages 187–208. Springer.

Joseph L Fleiss. 1971. Measuring nominal scale agreement among many raters. *Psychological bulletin*, 76(5):378.

Jena D Hwang, Chandra Bhagavatula, Ronan Le Bras, Jeff Da, Keisuke Sakaguchi, Antoine Bosselut, and Yejin Choi. 2021. (comet-) atomic 2020: On symbolic and neural commonsense knowledge graphs. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, pages 6384–6392.

Liwei Jiang, Antoine Bosselut, Chandra Bhagavatula, and Yejin Choi. 2021. "I'm not mad": Commonsense implications of negation and contradiction. In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 4380–4397, Online. Association for Computational Linguistics.

Minju Kim, Beong-woo Kwak, Youngwook Kim, Hong-in Lee, Seung-won Hwang, and Jinyoung Yeo. 2022. Dual task framework for improving persona-grounded dialogue dataset. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 36, pages 10912–10920.

Diederik P Kingma and Jimmy Ba. 2014. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*.

Jiwei Li, Michel Galley, Chris Brockett, Jianfeng Gao, and Bill Dolan. 2016a. A diversity-promoting objective function for neural conversation models. In *Proceedings of the 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 110–119, San Diego, California. Association for Computational Linguistics.

Jiwei Li, Michel Galley, Chris Brockett, Georgios Spithourakis, Jianfeng Gao, and Bill Dolan. 2016b. A persona-based neural conversation model. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 994–1003, Berlin, Germany. Association for Computational Linguistics.

Margaret Li, Stephen Roller, Ilia Kulikov, Sean Welleck, Y-Lan Boureau, Kyunghyun Cho, and Jason Weston. 2020. Don't say that! making inconsistent dialogue unlikely with unlikelihood training. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 4715–4728, Online. Association for Computational Linguistics.

Qian Liu, Yihong Chen, Bei Chen, Jian-Guang Lou, Zixuan Chen, Bin Zhou, and Dongmei Zhang. 2020. You impress me: Dialogue generation via mutual persona perception. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 1417–1427, Online. Association for Computational Linguistics.

Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019. Roberta: A robustly optimized bert pretraining approach. *arXiv preprint arXiv:1907.11692*.

Andrea Madotto, Zhaojiang Lin, Chien-Sheng Wu, and Pascale Fung. 2019. Personalizing dialogue agents via meta-learning. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 5454–5459, Florence, Italy. Association for Computational Linguistics.

Bodhisattwa Prasad Majumder, Harsh Jhamtani, Taylor Berg-Kirkpatrick, and Julian McAuley. 2020. Like hiking? you probably enjoy nature: Persona-grounded dialog with commonsense expansions. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 9194–9206, Online. Association for Computational Linguistics.

Qiao Qian, Minlie Huang, Haizhou Zhao, Jingfang Xu, and Xiaoyan Zhu. 2018. Assigning personality/profile to a chatting machine for coherent conversation generation. In *Proceedings of the 27th International Joint Conference on Artificial Intelligence*, pages 4279–4285.

Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, Ilya Sutskever, et al. 2019. Language models are unsupervised multitask learners.

Maarten Sap, Ronan Le Bras, Emily Allaway, Chandra Bhagavatula, Nicholas Lourie, Hannah Rashkin, Brendan Roof, Noah A Smith, and Yejin Choi. 2019. Atomic: An atlas of machine commonsense for if-then reasoning. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, pages 3027–3035.

Haoyu Song, Yan Wang, Kaiyan Zhang, Wei-Nan Zhang, and Ting Liu. 2021. BoB: BERT over BERT for training persona-based dialogue models from limited personalized data. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 167–177, Online. Association for Computational Linguistics.

Haoyu Song, Yan Wang, Wei-Nan 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 the 58th Annual Meeting of the Association for Computational Linguistics*, pages 5821–5831, Online. Association for Computational Linguistics.

Robyn Speer, Joshua Chin, and Catherine Havasi. 2017. Conceptnet 5.5: An open multilingual graph of general knowledge. In *Proceedings of the AAAI conference on artificial intelligence*, volume 31.

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.

Sean Welleck, Ilia Kulikov, Stephen Roller, Emily Dinan, Kyunghyun Cho, and Jason Weston. 2019a. Neural text generation with unlikelihood training. *arXiv preprint arXiv:1908.04319*.

Sean Welleck, Jason Weston, Arthur Szlam, and Kyunghyun Cho. 2019b. Dialogue natural language inference. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 3731–3741, Florence, Italy. Association for Computational Linguistics.

Adina Williams, Nikita Nangia, and Samuel Bowman. 2018. A broad-coverage challenge corpus for sentence understanding through inference. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pages 1112–1122, New Orleans, Louisiana. Association for Computational Linguistics.

Thomas Wolf, Victor Sanh, Julien Chaumond, and Clement Delangue. 2019. Transfertransfo: A transfer learning approach for neural network based conversational agents. *arXiv preprint arXiv:1901.08149*.

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, Melbourne, Australia. Association for Computational Linguistics.

Tianyi Zhang, Varsha Kishore, Felix Wu, Kilian Q Weinberger, and Yoav Artzi. 2019. Bertscore: Evaluating text generation with bert. *arXiv preprint arXiv:1904.09675*.

A Implementation Details

A.1 Persona Expansion

Our persona expansion model was implemented with HuggingFace Transformers.[2] The encoder and decoder initialize from the publicly available BERT-base-uncased[3] and GPT-2-small[4] model, respectively. Both models are with 12 layers and 768 hidden sizes. In addition, we use an Adam optimizer (Kingma and Ba, 2014), and the learning rate varies from 5e-6 to 5e-5. The α and β, hyperparameters for weighted NLL and UL training, are 0.8 and 0.07, respectively. The training was conducted on an Nvidia RTX3090 24G GPU with a batch size of 16. The best-performing model was found in epochs 11 to 13.

A.2 Response Generation

We utilized the repositories and implementations of GPT-2[5] and COMPAC[6] for response generation. To suit the purpose of our experiment, we adjusted some of the details of these models and trained them in a single-turn dialogue setting.

	Positive	Negative	All
# of unique anchor personas	4,253	12,540	16,793
# of silver profiles	2,428,926	3,710,299	6,139,225
Average # of profiles per unique anchor persona	571.11	295.88	365.58

Table 6: Statistics of silver profile dataset

Positive Profile				
Relation	xAttr	xNeed	xReact	xWant
Average # of expandable personas	4.48	3.65	2.60	4.55
Negative Profile				
Relation	xAttr	xNeed	xReact	xWant
Average # of expandable personas	3.17	1.24	2.74	3.50

Table 7: Expandable personas per relation

B Statistics of Silver Profile

The statistics of the silver profile dataset are shown in Table 6. We built a total of 6,139,225 synthetic profiles based on 16,793 unique anchor personas. Each profile has expandable personas for four relations. Table 7 shows the average number of expandable personas per relation.

[2] https://github.com/huggingface/transformers
[3] https://huggingface.co/bert-base-uncased
[4] https://huggingface.co/gpt2
[5] https://github.com/huggingface/transfer-learning-conv-ai
[6] https://github.com/majumderb/compac