KnowPAML: A Knowledge Enhanced Framework for Adaptable Personalized Dialogue Generation Using Meta-Learning

Aditya Shukla[§*]
TCS Research, India
adityashukla5597@gmail.com

Zishan Ahmad[§]
IIT Patna
1821cs18@iitp.ac.in

Asif Ekbal
IIT Patna
asif@iitp.ac.in

Abstract

In order to provide personalized interactions in a conversational system, responses must be consistent with the user and agent persona while still being relevant to the context of the conversation. Existing personalized conversational systems increase the consistency of the generated response by leveraging persona descriptions, which sometimes tend to generate irrelevant responses to the context. To solve this problems, we propose to extend the persona-agnostic meta-learning (PAML) framework (Madotto et al., 2019) by adding knowledge from ConceptNet knowledge graph (Speer et al.) with multi-hop attention mechanism (Tran and Niedereée, 2018). Knowledge is a concept in a triple form that helps in conversational flow. The multi-hop attention mechanism helps select the most appropriate triples with respect to the conversational context and persona description, as not all triples are beneficial for generating responses. The Meta-Learning (PAML) framework allows quick adaptation to different personas by utilizing only a few dialogue samples from the same user. Our experiments on the Persona-Chat dataset show that our method outperforms in terms of persona-adaptability, resulting in more persona-consistent responses, as evidenced by the entailment (Entl) score in the automatic evaluation and the consistency (Con) score in human evaluation.

1 Introduction

Recent advancements in personalized dialogue generation techniques that incorporate the personality of the speakers have enabled more human-like, natural, and persona-consistent responses. However, most methods require persona information in the form of style, persona profile, or persona statements, such as "I love meeting new people" and "Autumn is my favorite season" which can be very diverse and hence require a lot of data to model any persona type.

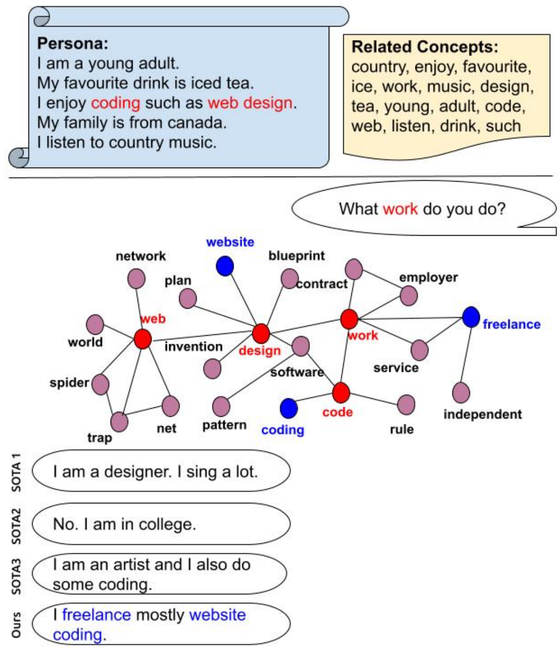

Figure 1: An example of developing a persona adaptable and knowledge guided response from test data using meta-learning and a commonsense knowledge graph. Concepts in red nodes come from the persona statement and dialogue history, whereas concepts in blue are in the generated response.

The Persona Agnostic Meta-Learning model (PAML) (Madotto et al., 2019) was developed to deal with these practical problems. This model, trained using meta-learning, would be able to adapt rapidly to new and unseen personas using only a few samples. The popular Model-Agnostic Meta-Learning (MAML) (Finn et al., 2017) framework served as the foundation for the PAML framework. Customized Model Agnostic Meta-Learning (CMAML) (Song et al., 2019) framework and Generating Personalized Dialogue via Multi-Task Meta-Learning (Lee et al., 2021) both largely follow the PAML framework except for an extra network structure optimization component and persona reconstruction component respectively. The Per-

[*]Work done at IIT Patna as part of M.Tech
[§]equal contribution

Proceedings of the 19th International Conference on Natural Language Processing (ICON), pages 194 - 203
December 15-18, 2022 ©2022 Association for Computational Linguistics

sonaChat corpus (Zhang et al., 2018), a popular personalized dialogue generating corpus that includes persona statements describing each speaker in addition to persona-specific dialogues, was used to benchmark all these frameworks. However, a comprehensive evaluation of these frameworks reveals that they continue to struggle to respond to situations and do not precisely match the persona statements. For example, in Figure 1, speaker1 asks a question about the nature of work that speaker2 is doing based on the persona statement: *I enjoy coding such as web design*. The responses generated by the State of the art frameworks are insufficiently precise. SOTA1, which is trained using dialogue history and persona statements and tested as a PAML framework, fails to accurately understand the persona. SOTA2, similar to SOTA1, is only trained on dialogue history and hence fails to generate the correct response. SOTA3 is a PAML framework that uses a meta-learning mechanism that wrongly understands the persona. Neither approach appears to be capable of accurately comprehending the persona.

Our model's goal is to appropriately interpret the context and persona statement while generating an engaging response that is both context and persona consistent. Real-life conversations, in general, begin with one topic and transition to other based on the personalities of the speakers. Basic commonplace knowledge also tends to spark conversations. Previously, external knowledge was used as a foundation for research, like "an open-domain knowledge graph" (Xing et al., 2017), "a commonsense knowledge base" (Zhou et al., 2018a), or "background documents" (Zhou et al., 2018b). Such external knowledge is incorporated into this research by enhancing the entity representations in dialogues with it and then generating responses based on it. This enhanced encoding of information in the model improves the quality of the generated responses. Also, commonsense knowledge-grounded (Majumder et al., 2020), knowledge expansion (Zhang et al., 2019), approaches were researched earlier, which also demonstrates an improvement in response generation.

In this work, we enhance the entity representation in persona-profile by using ConceptNet (Speer et al.) triples. We obtain triple representation using a graph encoding technique and incorporate these into the PAML framework (Madotto et al., 2019) using attention mechanism with the attention mechanism (Bahdanau et al., 2014). We then perform experiments on the PERSONA-CHAT dataset (Zhang et al., 2018) to test the effect of knowledge on persona adaptability. The PAML framework focuses on learning the different personas as independent tasks using the meta-learning approach, as opposed to building the model to represent all of the personas. The model is intended to be a few-shot learner using the PERSONA-CHAT where the train-test split contains a non-overlapping persona-type. The meta-learning training steps ensure that the model can swiftly adjust to a new unseen persona profile as well as the reaction style of a certain persona by utilizing only a few dialogues for training. Incorporating commonsense knowledge helps in a deeper understanding of the persona and dialogue context, therefore helping in adaptability.

The response generated by our model (c.f. Figure 1) is pretty much accurate and understands the persona statement correctly. Here, We extracted concepts from persona statement and dialogue context. For example Figure 1, "work" which is extracted from dialogue context is related to "freelance". Similarly, "web" (extracted from persona statement), "design" and "work" all three are related to "freelance" in some way, which is utilised to generate the response. Our experiment results show that our approach is effective in both automatic and human evaluation when compared to baseline models.

2 Related Work

Meta Learning The approach of teaching the model how to learn fast and effectively is known as meta-learning. Before, meta-learning was used in applications such as image classification. Finn et al. (2017) introduced MAML (Model-Agnostic Meta-Learning) technique, and it performed well for few-shot image classification. However, following MAML, various methodologies for NLP applications such as machine translation (Gu et al., 2018) and dialogue generation (Qian and Yu, 2019; Huang et al., 2020; Mi et al., 2019) were proposed, indicating an improvement. And following this, PAML (Madotto et al., 2019) was introduced, with a focus on personalized dialogue generation and Lee et al. (2021) proposed MTML and AMTML two frameworks in which they merged persona reconstruction task with the PAML which improved the persona consistency but failed in other automatic evaluation metrics.

Personalized dialogue generation have caught the interest of many in recent years, following Zhang et al. (2018)'s study of the task with the Persona-Chat dataset. Recent research focuses on advancing the dialogue generation by grounding persona information (Mazaré et al., 2018; Bao et al., 2019; Wolf et al., 2019) or bringing external knowledge such as the knowledge graphs (Long et al., 2017; Ghazvininejad et al., 2018) or supplementary texts (Vougiouklis et al., 2016; Xu et al., 2017) into the model. This research demonstrates that by doing so, the model becomes more informative and consistent with the personas of the speakers, hence improving generation performance.

Furthermore, if the knowledge graph is properly formed, or if it is domain-specific (Zhu et al., 2017; Xu et al., 2017), or if the knowledge base is large enough (Zhou et al., 2018a), the rich semantics representation is included through entities and relations (Hayashi et al., 2020). Based on this, Majumder et al. (2020) grounded the expanded persona statements using a commonsense knowledge graph, which aids in controlling the flow of the conversation. And In addition, reinforcement-learning-based framework was also proposed by Song et al. (2020) and Li et al. (2019) for making dialogue generation more informal. However, all of this work was based on either directly conditioning the response with the persona or by incorporating some commonsense knowledge into the model. None of the works attempted to generate a response from the persona except Madotto et al. (2019) and with the adaption of the knowledge graph.

3 Methodology

Our method extends the PAML framework (Madotto et al., 2019) by utilizing commonsense knowledge to generate personalized dialogues. The PAML is adapted from the MAML, which is capable of quickly adjusting to new, unknown tasks that were not employed during training. We continue to use the PERSONA-CHAT dataset (Zhang et al., 2018), which was used in the PAML framework. The dialogue in PERSONA-CHAT includes the utterances $u_{1:m}$ and persona statements $p_{1:n}$. In previous research, response $R = u_m$ was conditioned on the persona sentences $P = p_{1:n}$ and previous utterances $U = u_{1:m-1}$ following Equation 1:

$$f_W(R|U, P; W) = p(u_m|u_{1:m-1}, p_{1:n}; W) \quad (1)$$

Algorithm 1 KnowPAML

Require: P_m^{train}, P_m^{valid}
Require: Hyperparameters $\eta_{inner}, \eta_{outer}$
Require: iteration, patience, count
  Randomly initialize parameter W
  while $count < patience$ **do**
    persona batch from train set $P_{m_i}^t \sim P_m^{train}$
    for all $P_{m_i}^t$ **do**
      $(T_{p_i}, V_{p_i}) \sim P_{m_i}^t$
      calculate total loss
      $L_{T_{p_i}}^{total} = L_{T_{p_i}}^{g}(f_W) + L_{T_{p_i}}^{t}(f_W)$
      evaluate $\nabla_W L_{T_{p_i}}^{total}(f_W)$
      update $W_{p_i}' = W - \eta_{inner} \nabla_W L_{T_{p_i}}^{total}(f_W)$
    end for
    $W \leftarrow W - \eta_{outer} \nabla_W \frac{1}{S} \sum_{P_{m_i}} L_{V_{p_i}}(f_{W_{p_i}'})$
    persona batch from valid set
    $(T_{p_i'}, V_{p_i'}) \sim P_m^{valid}$
    if iteration % 10 == 0 **then**
      do **for** loop as above with $(T_{p_i'})$
      if $L_{V_{p_i}}(W') < L_{V_{p_i}}(W')_{best}$ **then**
        save weights
      else
        count+=1
      end if
    end if
  end while

In PAML, they first adapt W from the set of dialogue created by a persona and then respond conditioned only on the dialogue history rather than conditioning on both the dialogue history and persona sentences. So in this case Eq.(1) becomes:

$$f_W(R|U; W) = p(u_m|u_{1:m-1}; W) \quad (2)$$

3.1 Knowledgeable Persona-Agnostic Meta Learning

We use the PERSONA-CHAT dataset with ConceptNet triples associated with it at the utterance level (Section 3.2). First, we define P_m, which contains all of the personas and divide it into the train P_m^{train}, valid P_m^{valid}, and test P_m^{test} sets in the same way as PAML does. We sample persona batch $P_{m_i}^t$ from P_m^{train} for every training epoch, and then sample a set of utterances and associated triples as training T_{p_i} and another set as validating V_{p_i} from each persona in $P_{m_i}^t$. The utterance history in the batch is passed through a Transformer encoder and a representation $H = [h_1, h_2, ..., h_n]$

is obtained. Two weight matrices $W_{concept\_emb}$ and $W_{relation\_emb}$ are trained along with the model. The head and tail entity indexes are multiplied with $W_{concept\_emb}$ to obtain h_{emb} and t_{emb} respectively. The relation indexes of the triples are multiplied with $W_{relation\_emb}$ to obtain the representation r_{emb}. The final triple representation is obtained by concatenating these three representations $Trip = [h_{emb} : r_{emb} : t_{emb}]$. In this manner with each utterance we finally get a list of triple representations $T = [Trip_1, Trip_2, ..., Trip_n]$. Not all the triples in the list are useful for generating the appropriate knowledge-grounded response. Therefore we need to only select the triples appropriate with respect to the conversational context. To make this selection, we make use of the multi-hop attention mechanism (Tran and Niedereée, 2018). The attention mechanism works on a query q and an input sequence $T = [Trip_1, Trip_2, ..., Trip_n]$. For each k in K hop attention, the following steps are executed:

$$s_t^{(k)} = tanh(W_q^{(k)} Trip_t) \odot tanh(W_g^{(k)} g^{(k-1)}) \quad (3)$$

$$\alpha^{(k)} = softmax(w_s^{(k)^T} s_t^{(k)}) \quad (4)$$

$$o_q^{(k)} = \sum_t \alpha_t^{(k)} Trip_t \quad (5)$$

Here, $W_q^{(k)}$, $W_g^{(k)}$ and $w_s^{(k)}$ are the trainable parameters, and m is a separate memory vector for guiding the next attention step. It is recursively updated using the following equation:

$$g_q^{(k)} = g_q^{(k-1)} + o_q^k \odot q \quad (6)$$

The initial value of vector $g^{(0)}$ is defined based on the context vector $o_q^{(0)}$, given by the equation 7:

$$o_q^{(0)} = \frac{1}{l} \sum_t h_q(t) \odot q \quad (7)$$

The representation $o_q^{(k)}$ is the final attended and summed representation of T. We experiment with two settings for fusing this representation with the encoder hidden representation H: (i: AKnow-PAML). At each k the representation $o_q^{(k)}$ is added to each step of the encoded representation H to obtain the modified hidden representation H'. (ii: KnowPAML). At each k we concatenate o_q^k with h^k to obtain H''. We finally multiply this representation with W_{map} (a trainable matrix) to obtain the final modified hidden representation H'.

The decoder works on this obtained H' to produce the output. We use cross-entropy loss with decoder output and reference output to evaluate generation loss $L^g_{T_{p_i}}$ for dialogue model f_W, which is expressed as:

$$L^g_{T_{p_i}}(f_W) = -\sum logp(u_m | u_{1:m-1}, o_q^k; W) \quad (8)$$

Where, u_m is actual response, $u_{1:m-1}$ is dialogue history, o_q^k is attended representation of triples T and W is weight parameter.

In addition to generation loss, we evaluate the triple representation loss $L^t_{T_{p_i}}$ by using Equation 9 to train the optimal triple representation, following Bordes et al. (2013).

$$h_{emb} = t_{emb} - r_{emb} \quad (9)$$

During meta-training, total loss $L^{total}_{T_{p_i}}$ is calculated as the sum of generation loss and triple representation loss.

$$L^{total}_{T_{p_i}} = L^g_{T_{p_i}}(f_W) + L^t_{T_{p_i}}(f_W) \quad (10)$$

After training one batch of T_{p_i}, the model f_W, parameterized by W, is updated to $W\prime$ using SGD,

$$W'_{p_i} = W - \eta_{inner} \nabla_W L^{total}_{T_{p_i}}(f_W) \quad (11)$$

Where, η_{inner} is inner optimization learning rate and $L^{total}_{T_{p_i}}(f_W)$ is total training loss.

After the parameters are updated, meta-optimization is performed on the unseen dialogues from V_{p_i} set using the updated model $f(W'_{p_i})$ to enhance the model's performance. According Finn et al. (2017), the meta-objective is defined as follows:

$$\min_W \sum_{P_{m_i}} L_{V_{p_i}}(f_{W'_{p_i}}) = \sum_{P_{m_i}} L_{V_{p_i}}(f_{(W - \eta_{outer} \nabla_W L_{V_{p_i}}(f_W))}) \quad (12)$$

Where, $L_{V_{p_i}}(f_{W'_{p_i}})$ is the loss calculated on V_{p_i} set. During this step, we are only considering the dialogue generation loss rather than the total loss.

The initial parameters W are then adjusted using SGD by computing the gradient of average loss, which is the total of $L_{V_{p_i}}(f_{W'_{p_i}})$ obtained at each sampled persona divided by the batch size S. This is expressed as follows:

$$W \leftarrow W - \eta_{outer} \nabla_W \frac{1}{S} \sum_{P_{m_i}} L_{V_{p_i}}(f_{W'_{p_i}}) \quad (13)$$

Where, η_{outer} is the learning rate of outer optimization and $L_{V_{p_i}}(f_{W'_{p_i}})$ is the generation loss calculated on the V_{p_i} set by the updated model parameter $f_{W'_{p_i}}$. This is achieved by the use of second-order partial differentiation.

The validation set P_m^{valid} now validates this meta-training and meta-optimization after every ten iterations. We are dividing the training $T_{p'_i}$ and validating $V_{p'_i}$ set from the P_m^{valid} and performing the same step without altering the original parameters. And then save the best model based on the $L_{V_{p_i}}(W\prime)$. Also model get to know when to terminate training by increasing the count if the loss $L_{V_{p_i}}(W\prime)$ is not the best loss. Algorithm 1 gives an overview of the model.

3.2 Triple Retrieval

We use the ConceptNet, a commonsense knowledge graph, which connects words and phrases with labelled edges and has millions of concepts and edges associated with it. For every dialogue, we take concepts from the dialogue history of length one and associated persona statements to find the ConceptNet neighbours of each concept up to two hops. for example, zero-hop concept C_0 (which is taken from dialogue history and persona statements) is associated with one-hop concept C_1 (all immediate neighbours and all relation between them) and one-hop concept C_1 is associated with two-hop concept C_2. (head concept, relation, tail concept). The top 100 concepts-relations based on weights are then formed into triples (head concept, relation between them, tail concept).

4 Experiments

Our experiments are described in this section, including the dataset, implementation details, baselines, and evaluation metrics.

4.1 Dataset

The PERSONA-CHAT dataset (Zhang et al., 2018), which was also used in the PAML framework (Madotto et al., 2019), was employed for our experiments. The dataset has 1155 different personas in the train data and 100 each in the validation and test data. Each dialogue in this has 4 to 5 persona statements associated with it, and each unique persona has an average of 8.3 unique dialogues.

4.2 Implementation Details

We used the Transformer architecture (Vaswani et al., 2017) which includes six encoder, six decoder layers and four attention heads, just like Madotto et al. (2019), with Glove embedding (Pennington et al., 2014). Here, Transformer's hidden dimension and word embedding dimension are both set at 300. We utilized two different optimizers: SGD for training (inner loop optimizer) and ADAM (Kingma and Ba, 2014) for optimization (outer loop optimizer) with learning rates of 0.01 and 0.0003, respectively and the batch size is set to 16 for both the inner and outer loops.

4.3 Evaluation

We are employing both automatic and human evaluation metrics to evaluate the quality of response compared to baselines.

Automatic Evaluation We evaluate the perplexity (PPL), BLEU score (Papineni et al., 2002), and entailment score (entl), also known as the c score in Madotto et al. (2019). The PPL of the model indicates how well it understands the task on the test set; the lower the perplexity, the better the model understands the language. The BLEU score reflects how close the generated response is to the actual response; usually, a higher BLEU score suggests that the generated response is more comparable to the actual response; however, this cannot be asserted for every scenario. The entl score denotes how much persona information was included in the generated response from persona assertions. If the generated response entails the persona, the score is 1, if the response is independent of the persona, the score is 0, and if the response contradicts the persona information, the score is -1. The higher the entailment score, the more consistent the model is with the persona. This score is calculated using Madotto et al. (2019)'s fine-tuned BERT model, which was trained on persona-based Dialog NLI (Welleck et al., 2018) dataset and has an accuracy of 88.43%.

Human Evaluation We evaluate the Fluency (flcy) to measure the generated response's grammatical

Persona	
I go to the gym 4 days a week.	
I only drink water.	
I work in labor and delivery.	
I am happy being single and alone.	
I do not want children.	
Context	
Speaker 1:	hello. how are you doing ?
Speaker 2:	hi , i am doing great . how are you ?
Speaker 1:	feeling crabby , but i am like that naturally anyway .
Speaker 2:	oh ok . what do you do for a living ?
Speaker 1:	i am a kennel cleaner at a local animal shelter. and you ?
Speaker 2:	i am a doctor in the labor and deliver unit .
Speaker 1:	sounds very important . you must be a people person .
Speaker 2:	thanks i try to be . i love being along more than it looks like
Speaker 1:	i much prefer to hang with animals than people .
Speaker 2:	what do you do for fun ? i go work out 4 times a week
Speaker 1:	play video games and watch movies . you must be in good shape .
Speaker 2:	i try to be i will not drink soda or even tea , just water
Speaker 1:	that sounds like a healthy lifestyle .
Speaker 2:	it was hard to get use to at first . i use to love soda
Speaker 1:	i still love soda , especially sprite . do you have lots of friends or family ?
Responses	
Gold	i have family and a handful of friends when i am off i keep to myself
Corpus + Persona + Finetuning	i work at the hospital in labor and delivery
Corpus + Finetuning	i do not by the water i am okay and i am okay
PAML + Persona	i work at the gym at work at work
PAML	no it is easy for work
KnowPAML + Persona	i work at the hospital in labor and delivery
AKnowPAML	yes it does . it depends on the woman
KnowPAML	i do not have any , i like living alone .

Table 1: Example of responses generated by the implemented models using 10-shot.

correctness or readability, the Consistency (Con) to measure the persona information included in the response from persona assertions, and the Coherence (Coh) to measure the generated response's relevance with reference to the dialogue history. Three post-graduate-level human experts were asked to rate 200 randomly selected responses from 20 different personas which is generated by the proposed methodology. They were asked to rate consistency on a scale of -1 to 1, with -1 reflecting a contradiction of persona in the generated response, 0 reflecting neutral, and 1 reflecting persona-consistent. And fluency and coherence on a scale of 1 to 3. In fluency, 1 represents poor, 2 represents moderate, and 3 represents excellent in grammatical correctness or readability. And in coherence, 1 for inappropriate responses with context, 2 for moderately coherent responses with context, and 3 for contextually coherent responses.

4.4 Experimental Settings

In our research, we evaluated various training settings in the Transformer model:

PAML: Madotto et al. (2019) proposed this framework, It is a meta-trained model that is tested after fine-tuning the model with dialogues from the same persona.

PAML + Persona: Similar to PAML, except persona statements are included with dialogue history.

Corpus + Finetuning: A model is trained traditionally with dialogue history only as Eq.2 but fine-tuned and tested as PAML.

Experiment	10-Shot			5-Shot		
	PPL	BLEU	Entl	PPL	BLEU	Entl
Corpus + Persona + Finetuning	73.45	0.38	0.09	284.73	0.18	-0.05
Corpus (Varshney et al., 2020) + Finetuning	59.81	0.44	0.06	65.16	0.26	0.03
PAML + Persona	70.55	0.28	0.05	59.4	0.37	0.06
PAML (Madotto et al., 2019)	**41.88**	0.87	0.18	**40.55**	0.87	0.1
KnowPAML + Persona	56.32	0.87	0.11	55.66	0.8	0.06
AKnowPAML	45.78	**1.23**	0.18	48.6	**1.29**	0.08
KnowPAML	48.59	0.76	**0.24**	45.36	0.83	**0.15**

Table 2: Automatic Evaluation Results: When comparing with PAML, KnowPAML and AKnowPAML demonstrate that including concept into the model improves Persona Consistency and BLEU score respectively.

Experiment	10-Shot			5-Shot		
	flcy	Con	Coh	flcy	Con	Coh
Corpus + Persona + Finetuning	2.33	0.12	1.49	2.08	0.07	1.38
Corpus (Varshney et al., 2020) + Finetuning	2.03	0.08	1.53	1.91	0.06	1.46
PAML + Persona	2.36	0.09	1.35	2.11	0.05	1.29
PAML (Madotto et al., 2019)	2.77	0.24	2.17	2.63	0.13	2.06
KnowPAML + Persona	2.71	0.16	2.06	2.56	0.09	1.97
AKnowPAML	**2.89**	0.25	**2.33**	**2.74**	0.14	**2.23**
KnowPAML	2.83	**0.33**	2.21	2.64	**0.18**	2.13

Table 3: Human Evaluation Results: When comparing with PAML, KnowPAML and AKnowPAML demonstrate that including concept into the model improves Consistency (Con) and Coherence (Coh) respectively.

Corpus +Persona + Finetuning: Similar to the last one, except persona statements are included with dialogue history.

KnowPAML: We include knowledge in the form of triples with a multi-hop attention mechanism into the PAML using (setting (ii) section 3.1)

KnowPAML + Persona: Similar to KnowPAML, except persona statements are included with dialogue history.

AKnowPAML: This is similar to KnowPAML, but it differs in how attended triples are integrated into encoder outputs (setting (i) section 3.1).

We created dialogue history by appending all past utterances in the dialogue, with the length dictated by the number of turns that occurred, so there is no predetermined length. It will change with the turns, and there are no limits to the number of turns. We showed the results of 5 and 10 shots, where the response from the dialogue context and associated persona is generated after finetuning on 5 and 10 dialogue with the same persona, respectively.

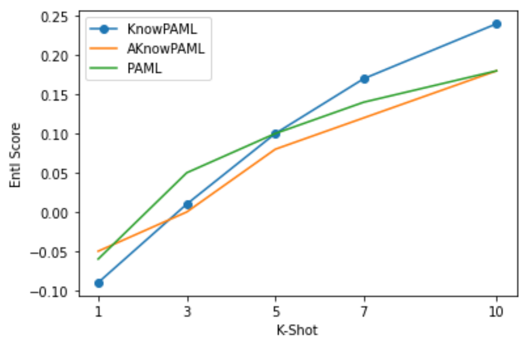

Figure 2: Entl Score vs K-Shot results of KnowPAML, AKnowPAML and PAML frameworks.

5 Result and Discussion

The results of the automatic and human evaluation for the 5-shot and 10-shot settings are shown in Table 2 and Table 3 respectively. KnowPAML was found to be most consistent to the persona in both automatic and human evaluation when compared to other systems. This implies that by incorporating persona and dialogue history knowledge into the existing framework, the system generates responses with a greater amount of persona information. AKnowPAML, on the other hand, performed better in terms of BLEU, fluency (flcy), and coherence (Coh). Perplexity (PPL), however, is higher in both instances, demonstrating a negative correlation between PPL and human likeness (Doğruöz and Skantze, 2021). KnowPAML outperformed PAML by 33.3% in 10-shots and 50% in 5-shots in terms of Entl score. In terms of consistency score (human evaluation), KnowPAML outperformed PAML by 37.5% in 10-shots and 38.5% in 5-shots. In the case of AKnowPAML, the Entl score is similar to the PAML score, while the BLEU score outperforms by 41.4% in 10-shots and 48.3% in 5-shots. Fluency and coherence are equivalent in PMAL, AKnowPAML, and KnowPAML in terms of human evaluation. The Entl score versus K-Shot results of KnowPAML, AKnowPAML, and PAML frameworks are shown in Figure 2, where K is 1, 3, 5, 7, and 10. Shots refer to the number of dialogues, such as 1-shot means one dialogue, 3-shot means three dialogue, that is used to fine-tune. Figure 2 shows that KnowPAML behaves more linearly compare to PAML and AKnowPAML. Entl score of KnowPAML intersects with PAML's score at 5-shot and then it keeps improving over PAML.

Persona:
i have got two more years in college i study law i want to be successful i am a student i have no siblings
Context:
hello , how are you this evening ?
Responses:
Corpus + Persona + Finetuning: i am well just studying
PAML + Persona: i am doing good and you ?
PAML: i am well just studying
KnowPAML + Persona: i am well just studying for class
AKnowPAML: fine . are you from one of my law classes ?
KnowPAML: fine . are you from one of my law classes ?

Table 4: Comparison of response generated by the PAML and KnowPAML using 10-shots.

5.1 Analysis

Table 1 shows an example of generated responses from different models using 10-shot fine tuning on held-out persona. As the PERSONA-CHAT dataset is an open domain conversation, the flow of the conversation can go in any direction, and the speaker generally talks in consonance with their persona. For new persona (not present in training data) it is often difficult to generate meaningful persona grounded responses that are also contextually relevant.

In Table 1, speaker 1 wants to know whether speaker 2 has a lot of friends or family; the response given by the KnowPAML framework is correct and consistent with the persona and context. In contrast, the response generated by the other framework is consistent with the persona but not with the context. The response from the PAML system in this case is inconsistent with both the persona and the context. In Table 4, The speaker asks, "How are you this evening?". All frameworks generate correct response given the context, however, our model generates a more engaging response that is consistent with one of the persona statements.

It has also been observed that responses generated by KnowPAML have an advantage in generating more persona consistent responses compared to AknowPAML, whereas AknowPAML attempts to generate depending on the real response, which is why the BLEU score is high. However, in terms of fluency, all are comparable, demonstrating that BLEU does not correlate with human judgment, as underlined by Liu et al. (2016) in their research of "How Not to Evaluate Your Dialogue System".

6 Conclusion

In this paper, we proposed the KnowPAML framework, which uses a multi-hop attention mechanism to absorb concepts in the form of triples from the ConceptNet knowledge graph in a Meta-learning setting. Our PERSONA-CHAT experiment demonstrates the advantage of using KnowPAML over previous frameworks in terms of persona-adaptability, resulting in more persona-consistent responses. The analysis of the generated responses reveals that the knowledge added model can successfully aid in persona adaptability, consistent response generation, and conversation flow. Although fluency and coherence are comparable to those of others, they can be improved further in the future by using a pre-trained language model.

Acknowledgement

The research reported in this paper is an outcome of the project “Autonomous Goal-Oriented and Knowledge-Driven Neural Conversational Agents", sponsored by Accenture LLP. Asif Ekbal acknowledges Visvesvaraya Young Faculty Research Fellowship (YFRF).

References

Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio. 2014. Neural machine translation by jointly learning to align and translate. *arXiv preprint arXiv:1409.0473*.

Siqi Bao, Huang He, Fan Wang, Rongzhong Lian, and Hua Wu. 2019. Know more about each other: Evolving dialogue strategy via compound assessment. *arXiv preprint arXiv:1906.00549*.

Antoine Bordes, Nicolas Usunier, Alberto Garcia-Duran, Jason Weston, and Oksana Yakhnenko. 2013. Translating embeddings for modeling multi-relational data. *Advances in neural information processing systems*, 26.

A Seza Doğruöz and Gabriel Skantze. 2021. How “open” are the conversations with open-domain chatbots? a proposal for speech event based evaluation. In *The 22nd Annual Meeting of the Special Interest Group on Discourse and Dialogue (SIGDIAL’22)*.

Chelsea Finn, Pieter Abbeel, and Sergey Levine. 2017. Model-agnostic meta-learning for fast adaptation of deep networks. In *International conference on machine learning*, pages 1126–1135. PMLR.

Marjan Ghazvininejad, Chris Brockett, Ming-Wei Chang, Bill Dolan, Jianfeng Gao, Wen-tau Yih, and Michel Galley. 2018. A knowledge-grounded neural conversation model. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 32.

Jiatao Gu, Yong Wang, Yun Chen, Kyunghyun Cho, and Victor OK Li. 2018. Meta-learning for low-resource neural machine translation. *arXiv preprint arXiv:1808.08437*.

Hiroaki Hayashi, Zecong Hu, Chenyan Xiong, and Graham Neubig. 2020. Latent relation language models. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, pages 7911–7918.

Yi Huang, Junlan Feng, Shuo Ma, Xiaoyu Du, and Xiaoting Wu. 2020. Towards low-resource semi-supervised dialogue generation with meta-learning. In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 4123–4128.

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

Jing Yang Lee, Kong Aik Lee, and Woon Seng Gan. 2021. Generating personalized dialogue via multi-task meta-learning. *arXiv preprint arXiv:2108.03377*.

Margaret Li, Stephen Roller, Ilia Kulikov, Sean Welleck, Y-Lan Boureau, Kyunghyun Cho, and Jason Weston. 2019. Don’t say that! making inconsistent dialogue unlikely with unlikelihood training. *arXiv preprint arXiv:1911.03860*.

Chia-Wei Liu, Ryan Lowe, Iulian V Serban, Michael Noseworthy, Laurent Charlin, and Joelle Pineau. 2016. How not to evaluate your dialogue system: An empirical study of unsupervised evaluation metrics for dialogue response generation. *arXiv preprint arXiv:1603.08023*.

Yinong Long, Jianan Wang, Zhen Xu, Zongsheng Wang, Baoxun Wang, and Zhuoran Wang. 2017. A knowledge enhanced generative conversational service agent. In *Proceedings of the 6th Dialog System Technology Challenges (DSTC6) Workshop*.

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.

Bodhisattwa Prasad Majumder, Harsh Jhamtani, Taylor Berg-Kirkpatrick, and Julian McAuley. 2020. Like hiking? you probably enjoy nature: Persona-grounded dialog with commonsense expansions. *arXiv preprint arXiv:2010.03205*.

Pierre-Emmanuel Mazaré, Samuel Humeau, Martin Raison, and Antoine Bordes. 2018. Training millions of personalized dialogue agents. *arXiv preprint arXiv:1809.01984*.

Fei Mi, Minlie Huang, Jiyong Zhang, and Boi Faltings. 2019. Meta-learning for low-resource natural language generation in task-oriented dialogue systems. *arXiv preprint arXiv:1905.05644*.

Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. 2002. Bleu: a method for automatic evaluation of machine translation. In *Proceedings of the 40th annual meeting of the Association for Computational Linguistics*, pages 311–318.

Jeffrey Pennington, Richard Socher, and Christopher D Manning. 2014. Glove: Global vectors for word representation. In *Proceedings of the 2014 conference on empirical methods in natural language processing (EMNLP)*, pages 1532–1543.

Kun Qian and Zhou Yu. 2019. Domain adaptive dialog generation via meta learning. *arXiv preprint arXiv:1906.03520*.

Haoyu Song, Wei-Nan Zhang, Jingwen Hu, and Ting Liu. 2020. Generating persona consistent dialogues by exploiting natural language inference. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, pages 8878–8885.

Yiping Song, Zequn Liu, Wei Bi, Rui Yan, and Ming Zhang. 2019. Learning to customize model structures for few-shot dialogue generation tasks. *arXiv preprint arXiv:1910.14326*.

R Speer, J Chin, and C ConceptNet Havasi. 5.5: An open multilingual graph of general knowledge. In *Proceedings of the Thirty-First AAAI Conference on Artificial Intelligence (December 2016)*, pages 4444–4451.

Nam Khanh Tran and Claudia Niedereée. 2018. Multi-hop attention networks for question answer matching. In *The 41st International ACM SIGIR Conference on Research & Development in Information Retrieval*, pages 325–334.

Deeksha Varshney, Asif Ekbal, Ganesh Prasad Nagaraja, Mrigank Tiwari, Abhijith Athreya Mysore Gopinath, and Pushpak Bhattacharyya. 2020. Natural language generation using transformer network in an open-domain setting. In *International Conference on Applications of Natural Language to Information Systems*, pages 82–93. Springer.

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.

Pavlos Vougiouklis, Jonathon Hare, and Elena Simperl. 2016. A neural network approach for knowledge-driven response generation. In *Proceedings of COLING 2016, the 26th International Conference on Computational Linguistics: Technical Papers*, pages 3370–3380.

Sean Welleck, Jason Weston, Arthur Szlam, and Kyunghyun Cho. 2018. Dialogue natural language inference. *arXiv preprint arXiv:1811.00671*.

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*.

Chen Xing, Wei Wu, Yu Wu, Jie Liu, Yalou Huang, Ming Zhou, and Wei-Ying Ma. 2017. Topic aware neural response generation. In *Proceedings of the AAAI conference on artificial intelligence*, volume 31.

Zhen Xu, Bingquan Liu, Baoxun Wang, Chengjie Sun, and Xiaolong Wang. 2017. Incorporating loose-structured knowledge into conversation modeling via recall-gate lstm. In *2017 international joint conference on neural networks (IJCNN)*, pages 3506–3513. IEEE.

Houyu Zhang, Zhenghao Liu, Chenyan Xiong, and Zhiyuan Liu. 2019. Grounded conversation generation as guided traverses in commonsense knowledge graphs. *arXiv preprint arXiv:1911.02707*.

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? *arXiv preprint arXiv:1801.07243*.

Hao Zhou, Tom Young, Minlie Huang, Haizhou Zhao, Jingfang Xu, and Xiaoyan Zhu. 2018a. Commonsense knowledge aware conversation generation with graph attention. In *IJCAI*, pages 4623–4629.

Kangyan Zhou, Shrimai Prabhumoye, and Alan W Black. 2018b. A dataset for document grounded conversations. *arXiv preprint arXiv:1809.07358*.

Wenya Zhu, Kaixiang Mo, Yu Zhang, Zhangbin Zhu, Xuezheng Peng, and Qiang Yang. 2017. Flexible end-to-end dialogue system for knowledge grounded conversation. *arXiv preprint arXiv:1709.04264*.