

Personalizing Dialogue Agents via Meta-Learning

Andrea Madotto,[†] Zhaojiang Lin[†], Chien-Sheng Wu, Pascale Fung

Center for Artificial Intelligence Research (CAiRE)

Department of Electronic and Computer Engineering

The Hong Kong University of Science and Technology, Clear Water Bay, Hong Kong

[amadotto, zlinao, cwuak, pascale]@ust.hk

Abstract

Existing personalized dialogue models use human designed persona descriptions to improve dialogue consistency. Collecting such descriptions from existing dialogues is expensive and requires hand-crafted feature designs. In this paper, we propose to extend Model-Agnostic Meta-Learning (MAML) (Finn et al., 2017) to personalized dialogue learning without using any persona descriptions. Our model learns to quickly adapt to new personas by leveraging only a few dialogue samples collected from the same user, which is fundamentally different from conditioning the response on the persona descriptions. Empirical results on Persona-chat dataset (Zhang et al., 2018) indicate that our solution outperforms non-meta-learning baselines using automatic evaluation metrics, and in terms of human-evaluated fluency and consistency.

1 Introduction

There is a growing interest in learning personalized chit-chat dialogue agents for making chatbots more consistent. Recently, a multi-turn conversational dataset called Persona-chat (Zhang et al., 2018) has been released, where two speakers are paired and a persona description (4-5 sentences) is randomly assigned to each of them. For example, “*I am an old man*” and “*I like to play football*” are one of the possible persona descriptions provided to the speaker. By conditioning the response generation on the persona descriptions, a chit-chat model is able to produce a more persona consistent dialogue (Zhang et al., 2018).

However, it is difficult to capture a persona just by using few sentences, and collecting a non-synthetic set of persona descriptions from a real human-human conversation, e.g., Reddit, is challenging as well since it requires hand-crafted fea-




Figure 1: The difference between finetuning from a) joint training on all personas and b) meta-learning persona. The solid line represents the optimization path of the initial parameters and dashed line the fine-tuning path. Meta-learned initial parameters can faster adapt to a new persona.

ture designs (Mazare et al., 2018). In light of this, we propose to leverage a set of dialogues done by the same persona directly, instead of using its persona descriptions, to generate a more consistent response.

We consider learning different personas as different tasks via meta-learning algorithms, which is fundamentally different from optimizing the model to represent all the personas. A high-level intuition of the difference between these two approaches is shown in Figure 1. We aim to learn a persona-independent model that is able to quickly adapt to a new persona given the dialogues. We formulate this task as a few-shot learning problem, where K dialogues are used for training and the remaining for the test. Hence, we expect to learn initial parameters of a dialogue model that can quickly adapt to the response style of a certain persona just by using few dialogues.

The main contribution of this paper is to cast the personalized dialogue learning as a meta-learning problem, which allows our model to generate personalized responses by efficiently leveraging only a few dialogue samples instead of human-designed

[†] These two authors contributed equally.

persona descriptions. Empirical results show that our solution outperforms joint training, in terms of human-evaluated fluency and consistency.

2 Personalized Dialogue Learning

2.1 Persona-conditioned dialogue

In Persona-chat dataset (Zhang et al., 2018), a dialogue is defined as a set of utterances $U = \{u_1, \dots, u_n\}$ and a persona description is defined as a set of sentences $P = \{p_1, \dots, p_m\}$. A personalized dialogue model f_θ is trained to produce a response $Y = u_t$ conditioned on previous utterances $X = \{u_1, \dots, u_{t-1}\}$ and persona sentences P :

$$f_\theta(Y|X, P; \theta) = p(u_t|u_{1:t-1}, p_{1:m}; \theta) \quad (1)$$

2.2 Persona-agnostic dialogue

Instead of conditioning our response on the persona sentences, we first adapt θ to the set of dialogue made by a persona P and then we only use the dialogue history to condition our response. Eq. (1) becomes:

$$f_\theta(Y|X; \theta) = p(u_t|u_{1:t-1}; \theta) \quad (2)$$

Therefore, we define the set of dialogues of a persona P as $\mathcal{D}_P = \{U_1, \dots, U_k\}$. Conceptually, a model f_θ is expected to generate personalized response after being trained with a few dialogues example from \mathcal{D}_P . The main idea of our work is to use Model-Agnostic Meta-Learning (MAML) (Finn et al., 2017) to learn an initial set of parameters that can quickly learn a persona from few dialogues sample. We refer to the proposed meta-learning method for persona dialogues as Persona-Agnostic Meta-Learning (PAML).

Persona-agnostic meta-learning (PAML) We define the persona meta-dataset as $\mathcal{D} = \{\mathcal{D}_{p_1}, \dots, \mathcal{D}_{p_z}\}$, where z is the number of persona. Before training, \mathcal{D} is split into $\mathcal{D}_{train}, \mathcal{D}_{valid}, \mathcal{D}_{test}$. For each training epoch, we uniformly sample a batch of personas \mathcal{D}_{p_i} from \mathcal{D}_{train} , then from each persona in \mathcal{D}_{p_i} we sample a set of dialogues as training $\mathcal{D}_{p_i}^{train}$, and another set of dialogues as validation $\mathcal{D}_{p_i}^{valid}$. After t iterations of training on $\mathcal{D}_{p_i}^{train}$, the dialogue model f_θ , parameterized by θ , is updated to θ'_{p_i} by standard gradient descent,

$$\theta'_{p_i} = \theta - \alpha \nabla_\theta \mathcal{L}_{\mathcal{D}_{p_i}^{train}}(f_\theta) \quad (3)$$

Algorithm 1 Persona-Agnostic Meta-Learning

Require: \mathcal{D}_{train}

Require: α, β : step size hyperparameters

```

1: Randomly initialize  $\theta$ 
2: while not done do
3:   Sample batch of persona  $\mathcal{D}_{p_i} \sim \mathcal{D}_{train}$ 
4:   for all  $\mathcal{D}_{p_i}$  do
5:      $(\mathcal{D}_{p_i}^{train}, \mathcal{D}_{p_i}^{valid}) \sim \mathcal{D}_{p_i}$ 
6:     Evaluate  $\nabla_\theta \mathcal{L}_{\mathcal{D}_{p_i}^{train}}(f_\theta)$  using  $\mathcal{D}_{p_i}^{train}$ 
7:     Compute adapted parameters with
       gradient descent:
        $\theta'_{p_i} = \theta - \alpha \nabla_\theta \mathcal{L}_{\mathcal{D}_{p_i}^{train}}(f_\theta)$ 
8:   end for
9:    $\theta \leftarrow \theta - \beta \sum_{\mathcal{D}_{p_i} \sim \mathcal{D}_{train}} \nabla_\theta \mathcal{L}_{\mathcal{D}_{p_i}^{valid}}(f_{\theta'_{p_i}})$ 
10: end while

```

where α is learning of the inner optimization, and $\mathcal{L}_{\mathcal{D}_{p_i}^{train}}$ the training loss. Specifically, cross-entropy loss is used for training the response generation:

$$\mathcal{L}_{\mathcal{D}_{p_i}}(f_\theta) = - \sum_{\mathcal{D}_{p_i}} \log p(u_t|u_{1:t-1}; \theta) \quad (4)$$

The meta-learning model is then trained to maximize the performance of the adapted model $f_{\theta'_{p_i}}$ to the unseen dialogues in $\mathcal{D}_{p_i}^{valid}$. Following Finn et al. (2017), we define the meta-objective as:

$$\min_{\theta} \sum_{\mathcal{D}_{p_i} \sim \mathcal{D}_{train}} \mathcal{L}_{\mathcal{D}_{p_i}^{valid}}(f_{\theta'_{p_i}}) = \sum_{\mathcal{D}_{p_i} \sim \mathcal{D}_{train}} \mathcal{L}_{\mathcal{D}_{p_i}^{valid}}(f_{\theta - \alpha \nabla_\theta \mathcal{L}_{\mathcal{D}_{p_i}^{train}}(f_\theta)}) \quad (5)$$

where $\mathcal{L}_{\mathcal{D}_{p_i}^{valid}}(f_{\theta'_{p_i}})$ is the loss evaluated on $\mathcal{D}_{p_i}^{valid}$. For optimizing Eq.(5), we apply again stochastic gradient descent on the meta-model parameters θ by computing the gradient of $\mathcal{L}_{\mathcal{D}_{p_i}^{valid}}(f_{\theta'_{p_i}})$, which is:

$$\theta \leftarrow \theta - \beta \sum_{\mathcal{D}_{p_i} \sim \mathcal{D}_{train}} \nabla_\theta \mathcal{L}_{\mathcal{D}_{p_i}^{valid}}(f_{\theta'_{p_i}}) \quad (6)$$

where β is meta-learning rate. This process requires second order optimization partial derivatives, which can be computed by any automatic differentiation library (e.g. PyTorch, Tensorflow etc.). A summary of the training procedure is shown in Algorithm 1.

	<i>Automatic</i>			<i>Human</i>	
	PPL	BLEU	C	Fluency	Consistency
<i>Human</i>	-	-	0.33	3.434	0.234
<i>Dialogue+Persona</i>	30.42	1.00	0.07	3.053	0.011
<i>Dialogue</i>	36.75	0.64	-0.03	-	-
<i>Dialogue+Fine-tuning</i>	32.96	0.90	0.00	3.103	0.038
<i>PAML</i>	41.64	0.74	0.20	3.185	0.197

Table 1: Results of automatic and human evaluation: *PAML* vs *Dialogue+Persona* shows the our approach can achieve good consistency by using few dialogues instead of conditioning on the persona description, *PAML* vs *Dialogue+Fine-tuning* shows the effectiveness of meta-learning approach in personalizing dialogue model.

3 Experiment and Results

The experiments are conducted using Persona-chat (Zhang et al., 2018). To create the meta-sets \mathcal{D} , we match the dialogues by their persona description separately for train, validation and test, by following the same persona split as in Zhang et al. (2018). On average each persona description has 8.3 unique dialogues. In the Appendix, we report the number of dialogue distribution.

Experimental setting In our experiments, we compared different training settings: (*Dialogue*) a model trained using dialogue history, as in Eq.(2); (*PAML*) a meta-trained model as in Eq.(5), where we test each set $\mathcal{D}_{p_i} \in \mathcal{D}_{test}$ by selecting one dialogue and training with all the others. To elaborate, suppose we are testing $U_t \in \mathcal{D}_{p_i}$ then we first fine-tuning using all the dialogues in $\mathcal{D}_{p_i} \setminus U_t$, and then test on U_t . This process is repeated for all the dialogues in \mathcal{D}_{p_i} . (*Dialogue+Fine-tuning*) we use the same testing as *PAML* but on a model trained as *Dialogue*. We also report a trained model that assumes persona description is available and we refer it as (*Dialogue+Persona*).

Implementation details We implemented f_θ using a standard Transformer architecture (Vaswani et al., 2017) with pre-trained Glove embedding (Pennington et al., 2014)¹. For the standard training, we used Adam (Kingma and Ba, 2014) optimizer with a warm-up learning rate strategy, and a batch size of 32. Instead, in meta-training, we used SGD for the inner loop and Adam for the outer loop with learning rate $\alpha = 0.01$ and $\beta = 0.0003$ respectively, and batch size of 16 for both. In all the model we used beam search with beam size 5.

¹The model and the pre-processing scripts are available at <https://github.com/HLTCHKUST/PAML>

3.1 Evaluation metric

The objective of the evaluation is to verify whether *PAML* can produce a more consistent response with reference to the given dialogue and persona description (even though is not seen). To do so, we employ both automatic and human evaluation.

Automatic We report perplexity and BLEU score (Papineni et al., 2002) of the generate sentences against the human-generated prediction. Aside of standards evaluation metrics, we also train a Natural Language Inference (NLI) model using Dialog NLI (Sean et al., 2018) dataset, a recently proposed corpus based on Persona dataset, with NLI annotation between persona description sentences and dialogues utterance. We fine-tune a pre-trained BERT model (Devlin et al., 2018) using the DNLI corpus and achieve a test set accuracy of 88.43%, which is aligned to the best-reported model ESIM (Chen et al., 2017) in Sean et al. (2018) (with 88.20% accuracy). Then, we defined a new evaluation metric for dialogue consistency as follow:

$$\text{NLI}(u, p_j) = \begin{cases} 1 & \text{if } u \text{ entails } p_j \\ 0 & \text{if } u \text{ is independent to } p_j \\ -1 & \text{if } u \text{ contradicts } p_j \end{cases}$$

$$\mathcal{C}(u) = \sum_j^m \text{NLI}(u, p_j) \quad (7)$$

where u is a generated utterance and the p_j is one sentence in the persona description. Hence, having a higher consistency \mathcal{C} score means having a more persona consistent dialogue response.

Human Since automatic evaluation performs poorly in this task (Liu et al., 2016), we perform a human evaluation using crowd-sourced workers. We randomly selected 300 generated response examples from 10 unique personas and we asked

each worker to evaluate fluency (1 to 5) and consistency of the generated response with respect to the dialogue history and the respective persona description. We asked the workers to assign a score of 1, 0 or -1 for consistent, neutral, and contradicts respectively, the full instruction set is available in the Appendix.

3.2 Results

Table 1 shows both automatic and human evaluation results. *PAML* achieve consistently better results in term of dialogue consistency in both automatic and human evaluation. The latter also shows that all the experimental settings have comparable fluency scores, where instead perplexity and BLEU score are lower in *PAML*. This confirms that these measures are not correlated to human judgment (Liu et al., 2016). For completeness, we also show generated responses examples from *PAML* and baseline models in Appendix.

On the other hand, the human evaluated consistency is aligned to the C score, which confirms the meaningfulness of the defined measure. This agrees with results of Sean et al. (2018), where the authors showed that by re-ranking the beam search hypothesis using the DNLI score (i.e. C score), they achieved a substantial improvement in dialogue consistency.

Few-shot Learning We analyze the ability of our model to fast adapt to a certain persona in term of shots. We define shot as the number of dialogues used in $\mathcal{D}_{p_i}^{train}$ for fine-tuning a certain persona, e.g. 1-shot one dialogue, 3-shot three dialogue and so on. Figure 2 compares the k -shot consistency C results for k equal to 0, 1, 3, 5 and 10, both *PAML* and *Dialogue+Fine-tuning*. *PAML* can achieve a high consistency score just by using 3 dialogues, which is better than *Persona+Dialogue*. On the other hand, *Dialogue+Fine-tuning* cannot properly leverage the dialogues in \mathcal{D}_{p_i} , which proves the effectiveness of training with meta-learning.

4 Related Work

Meta-Learning Meta-learning (Thrun and Pratt, 1998; Schmidhuber, 1987, 1992; Naik and Mammone, 1992; Bengio et al., 1992) is sub-field of machine learning with the aim of learning the learning algorithm itself. Recently, several meta-learning models has been proposed for solving few-shot image classification (Ravi




Figure 2: k-shot results for different settings. Consistency of *PAML* grows linearly with respect to k .

and Larochelle, 2016; Vinyals et al., 2016; Finn et al., 2017; Mishra et al., 2017; Santoro et al., 2016), optimization (Andrychowicz et al., 2016) and reinforcement learning (Finn et al., 2017). Meta-learning for NLP application is less common, and it has been applied in semantic parsing task (Huang et al., 2018), machine translation for low resource language (Gu et al., 2018), and for text classification (Yu et al., 2018). To the best of our knowledge, this is the first attempt in adapting meta-learning to personalized dialogue learning.

Personalized Dialogue Li et al. (2016) was the first to propose a persona based dialogue models for improving response consistency. Zhang et al. (2018) introduced Persona-chat, which was further extended in ConvAI2 (2019). Several works improved on the initial baselines with various methodologies (Kulikov et al., 2018; Yavuz et al.; Hancock et al., 2019; Lucas et al., 2009; Joshi et al., 2017; Zemlyanskiy and Sha, 2018; Gao et al., 2018). However, all of these previous works conditioned their response on the persona description, instead of using the dialogues produced by the persona.

5 Conclusion

In this paper, we present a novel meta-learning setting for personalizing dialogue agents without conditioning the model response to the persona description. This is especially useful since obtaining such persona description requires human effort. Moreover, we show that a dialogue agent trained with meta-learning achieves a more consistent dialogue by both of automatic measures and human evaluation. In future works, we plan to apply meta-learning to comment genera-

tion (Lin et al., 2019) and task-oriented dialogues systems (Madotto et al., 2018; Wu et al., 2019, 2017, 2018; Reddy et al., 2018).

6 Acknowledgments

This work has been funded by MRP/055/18 of the Innovation Technology Commission, of the Hong Kong University of Science and Technology.

References

- Marcin Andrychowicz, Misha Denil, Sergio Gomez, Matthew W Hoffman, David Pfau, Tom Schaul, Brendan Shillingford, and Nando De Freitas. 2016. Learning to learn by gradient descent by gradient descent. In *Advances in Neural Information Processing Systems*, pages 3981–3989.
- Samy Bengio, Yoshua Bengio, Jocelyn Cloutier, and Jan Gecsei. 1992. On the optimization of a synaptic learning rule. In *Preprints Conf. Optimality in Artificial and Biological Neural Networks*, pages 6–8. Univ. of Texas.
- Qian Chen, Xiaodan Zhu, Zhen-Hua Ling, Si Wei, Hui Jiang, and Diana Inkpen. 2017. [Enhanced lstm for natural language inference](#). In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1657–1668. Association for Computational Linguistics.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2018. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*.
- Emily Dinan, Varvara Logacheva, Valentin Malykh, Alexander Miller, Kurt Shuster, Jack Urbanek, Douwe Kiela, Arthur Szlam, Iulian Serban, Ryan Lowe, et al. 2019. The second conversational intelligence challenge (convai2). *arXiv preprint arXiv:1902.00098*.
- Chelsea Finn, Pieter Abbeel, and Sergey Levine. 2017. Model-agnostic meta-learning for fast adaptation of deep networks. In *Proceedings of the 34th International Conference on Machine Learning-Volume 70*, pages 1126–1135. JMLR. org.
- Jianfeng Gao, Michel Galley, and Lihong Li. 2018. Neural approaches to conversational ai. In *The 41st International ACM SIGIR Conference on Research & Development in Information Retrieval*, pages 1371–1374. ACM.
- Jiatao Gu, Yong Wang, Yun Chen, Victor O. K. Li, and Kyunghyun Cho. 2018. [Meta-learning for low-resource neural machine translation](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 3622–3631. Association for Computational Linguistics.
- Braden Hancock, Antoine Bordes, Pierre-Emmanuel Mazare, and Jason Weston. 2019. Learning from dialogue after deployment: Feed yourself, chatbot! *arXiv preprint arXiv:1901.05415*.
- Po-Sen Huang, Chenglong Wang, Rishabh Singh, Wen-tau Yih, and Xiaodong He. 2018. [Natural language to structured query generation via meta-learning](#). In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 2 (Short Papers)*, pages 732–738. Association for Computational Linguistics.
- Chaitanya K Joshi, Fei Mi, and Boi Faltings. 2017. Personalization in goal-oriented dialog. *arXiv preprint arXiv:1706.07503*.
- Diederik P Kingma and Jimmy Ba. 2014. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*.
- Ilya Kulikov, Alexander H Miller, Kyunghyun Cho, and Jason Weston. 2018. Importance of a search strategy in neural dialogue modelling. *arXiv preprint arXiv:1811.00907*.
- Jiwei Li, Michel Galley, Chris Brockett, Georgios Spithourakis, Jianfeng Gao, and Bill Dolan. 2016. A persona-based neural conversation model. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, volume 1, pages 994–1003.
- Zhaojiang Lin, Genta Indra Winata, and Pascale Fung. 2019. Learning comment generation by leveraging user-generated data. In *ICASSP 2019-2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 7225–7229. IEEE.
- Chia-Wei Liu, Ryan Lowe, Iulian Serban, Mike 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](#). In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 2122–2132. Association for Computational Linguistics.
- JM Lucas, F Fernández, J Salazar, J Ferreiros, and R San Segundo. 2009. Managing speaker identity and user profiles in a spoken dialogue system. *Procesamiento del lenguaje natural*, (43).
- Andrea Madotto, Chien-Sheng Wu, and Pascale Fung. 2018. Mem2seq: Effectively incorporating knowledge bases into end-to-end task-oriented dialog systems. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1468–1478.
- Pierre-Emmanuel Mazare, Samuel Humeau, Martin Raison, and Antoine Bordes. 2018. [Training millions of personalized dialogue agents](#). In *Proceedings of the 2018 Conference on Empirical Methods*

- in *Natural Language Processing*, pages 2775–2779. Association for Computational Linguistics.
- Nikhil Mishra, Mostafa Rohaninejad, Xi Chen, and Pieter Abbeel. 2017. A simple neural attentive meta-learner. *ICLR*.
- Devang K Naik and RJ Mammone. 1992. Meta-neural networks that learn by learning. In *[Proceedings 1992] IJCNN International Joint Conference on Neural Networks*, volume 1, pages 437–442. IEEE.
- 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 on association for computational linguistics*, pages 311–318. Association for Computational Linguistics.
- Jeffrey Pennington, Richard Socher, and Christopher 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.
- Sachin Ravi and Hugo Larochelle. 2016. Optimization as a model for few-shot learning.
- Revanth Reddy, Danish Contractor, Dinesh Raghu, and Sachindra Joshi. 2018. Multi-level memory for task oriented dialogs. *arXiv preprint arXiv:1810.10647*.
- Adam Santoro, Sergey Bartunov, Matthew Botvinick, Daan Wierstra, and Timothy Lillicrap. 2016. Meta-learning with memory-augmented neural networks. In *International conference on machine learning*, pages 1842–1850.
- Jurgen Schmidhuber. 1987. [Evolutionary principles in self-referential learning. on learning now to learn: The meta-meta-meta...-hook](#). Diploma thesis, Technische Universitat Munchen, Germany, 14 May.
- Jürgen Schmidhuber. 1992. Learning to control fast-weight memories: An alternative to dynamic recurrent networks. *Neural Computation*, 4(1):131–139.
- Welleck Sean, Jason Weston, Arthur Szlam, and Kyunghyun Cho. 2018. Dialogue natural language inference. *arXiv preprint arXiv:1811.00671*.
- Sebastian Thrun and Lorien Pratt, editors. 1998. *Learning to Learn*. Kluwer Academic Publishers, Norwell, MA, USA.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. [Attention is all you need](#). In I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett, editors, *Advances in Neural Information Processing Systems 30*, pages 5998–6008. Curran Associates, Inc.
- Oriol Vinyals, Charles Blundell, Timothy Lillicrap, Daan Wierstra, et al. 2016. Matching networks for one shot learning. In *Advances in neural information processing systems*, pages 3630–3638.
- Chien-Sheng Wu, Andrea Madotto, Genta Winata, and Pascale Fung. 2017. End-to-end recurrent entity network for entity-value independent goal-oriented dialog learning. In *Dialog System Technology Challenges Workshop, DSTC6*.
- Chien-Sheng Wu, Andrea Madotto, Genta Winata, and Pascale Fung. 2018. End-to-end dynamic query memory network for entity-value independent task-oriented dialog. In *2018 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 6154–6158. IEEE.
- Chien-Sheng Wu, Richard Socher, and Caiming Xiong. 2019. [Global-to-local memory pointer networks for task-oriented dialogue](#). In *International Conference on Learning Representations*.
- Semih Yavuz, Abhinav Rastogi, Guan-lin Chao, Dilek Hakkani-Tür, and Amazon Alexa AI. Deepcopy: Grounded response generation with hierarchical pointer networks.
- Mo Yu, Xiaoxiao Guo, Jinfeng Yi, Shiyu Chang, Saloni Potdar, Yu Cheng, Gerald Tesauro, Haoyu Wang, and Bowen Zhou. 2018. [Diverse few-shot text classification with multiple metrics](#). 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 1206–1215. Association for Computational Linguistics.
- Yury Zemlyanskiy and Fei Sha. 2018. Aiming to know you better perhaps makes me a more engaging dialogue partner. *CoNLL 2018*, page 551.
- 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. Association for Computational Linguistics.