

Few-Shot Dialogue Generation Without Annotated Data: A Transfer Learning Approach

Igor Shalyminov[†], Sungjin Lee[‡], Arash Eshghi[†], and Oliver Lemon[†]

[†]Heriot-Watt University, UK

[‡]Microsoft Research, US

[†]{is33, a.eshghi, o.lemon}@hw.ac.uk, [‡]sungjinlee.plus@gmail.com

Abstract

Learning with minimal data is one of the key challenges in the development of practical, production-ready goal-oriented dialogue systems. In a real-world enterprise setting where dialogue systems are developed rapidly and are expected to work robustly for an ever-growing variety of domains, products, and scenarios, efficient learning from a limited number of examples becomes indispensable.

In this paper, we introduce a technique to achieve state-of-the-art dialogue generation performance in a few-shot setup, without using any annotated data. We do this by leveraging background knowledge from a larger, more highly represented dialogue source — namely, the MetaLWOz dataset. We evaluate our model on the Stanford Multi-Domain Dialogue Dataset, consisting of human-human goal-oriented dialogues in in-car navigation, appointment scheduling, and weather information domains.

We show that our few-shot approach achieves state-of-the-art results on that dataset by consistently outperforming the previous best model in terms of BLEU and Entity F1 scores, while being more data-efficient by not requiring any data annotation.

1 Introduction

Data-driven dialogue systems are becoming widely adopted in enterprise environments. One of the key properties of a dialogue model in this setting is its *data efficiency*, i.e. whether it can attain high accuracy and good generalization properties when only trained from minimal data.

Recent deep learning-based approaches to training dialogue systems (Ultes et al., 2018; Wen et al., 2017) put emphasis on collecting large amounts of data in order to account for numerous variations in the user inputs and to cover as many

dialogue trajectories as possible. However, in real-world production environments there isn't enough domain-specific data easily available throughout the development process. In addition, it's important to be able to rapidly adjust a system's behavior according to updates in requirements and new product features in the domain. Therefore, data-efficient training is a priority direction in dialogue system research.

In this paper, we build on a technique to train a dialogue model for a new domain in a 'zero-shot' setup (in terms of full dialogues in the target domain) only using *annotated* 'seed' utterances (Zhao and Eskénazi, 2018).

We present an alternative, 'few-shot' approach to data-efficient dialogue system training: we do use complete in-domain dialogues while using approximately the same amount of training data as Zhao and Eskénazi (2018), with respect to utterances. However, in our method, *no annotation is required* — we instead use a latent dialogue act annotation learned in an unsupervised way from a larger (multi-domain) data source, broadly following the model of Zhao et al. (2018). This approach is potentially more attractive for practical purposes because it is easier to collect unannotated dialogues than collecting utterances across various domains under a consistent annotation scheme.

2 Related Work

There is a substantial amount of work on learning dialogue with minimal data — starting with the Dialog State Tracking Challenge 3 (Henderson et al., 2014) where the problem was to adjust a pre-trained state tracker to a different domain using a seed dataset.

In dialogue response generation, there has also been work on bootstrapping a goal-oriented dialogue system from a few examples using a lin-

guistically informed model: (Eshghi et al., 2017) used an incremental semantic parser – DyLan (Eshghi et al., 2011; Eshghi, 2015) – to obtain contextual meaning representations, and based the dialogue state on this (Kalatzis et al., 2016). Incremental response generation was learned using Reinforcement Learning, again using the parser to incrementally process the agent’s output and thus prune ungrammatical paths for the learner. Compared to a neural model — End-to-End Memory Network (Sukhbaatar et al., 2015), this linguistically informed model was superior in a 1-shot setting (Shalymov et al., 2017). At the same time, its main linguistic resource — a domain-general dialogue grammar for English — makes the model inflexible unless wide coverage is achieved.

Transfer learning for Natural Language Processing is strongly motivated by recent advances in vision. When training a convolutional neural network (CNN) on a small dataset for a specific problem domain, it often helps to learn low-level convolutional features from a greater, more diverse dataset. For numerous applications in vision, ImageNet (Deng et al., 2009) became the source dataset for pre-training convolutional models. For NLP, the main means for transfer were Word2Vec word embeddings (Mikolov et al., 2013) which have recently been updated to models capturing contexts as well (Peters et al., 2018; Devlin et al., 2018). While these tools are widely known to improve performance in various tasks, more specialized models could as well be created for specific research areas, e.g. dialogue generation in our case.

The models above are some of the approaches to one of the central issues of efficient knowledge transfer — learning a unified data representation generalizable across datasets, dubbed ‘representation learning’. In our approach, we will use one such technique based on variational autoencoding with discrete latent variables (Zhao et al., 2018). In this paper we present an approach to transfer learning which is more tailored — both model-wise and dataset-wise — to goal-oriented dialogue in underrepresented domains.

3 The approach

3.1 Zero-shot theoretical framework

We first describe the original Zero-Shot Dialogue Generation (ZSDG) theoretical framework of (Zhao and Eskénazi, 2018) which we base our

work on. For ZSDG, there is a set of source dialogue domains and one target domain, with the task of training a dialogue response generation model from all the available source data and a significantly reduced subset of the target data (referred to as *seed* data). The trained system’s performance is evaluated exclusively on the target domain.

More specifically, the data in ZSDG is organized as follows. There are unannotated dialogues in the form of $\{c, \mathbf{x}, d\}_{src/tgt}$ — tuples of dialogue contexts, responses, and domain names respectively for each of the source and target domains. There are also domain descriptions in the form of $\{\mathbf{x}, \mathbf{a}, d\}_{src/tgt}$ — tuples of utterances, slot-value annotations, and domain names respectively for source and target domains.

ZSDG is essentially a hierarchical encoder-decoder model which is trained in a multi-task fashion by receiving two types of data: (1) dialogue batches drawn from all the available source-domain data, and (2) seed data batches, a limited number of which are drawn from domain description data for all of the source and target domains.

ZSDG model optimizes for 2 objectives. With dialogue batches, the model maximizes the probability of generating a response given the context:

$$\mathcal{L}_{dialog} = -\log p_{\mathcal{F}^d}(\mathbf{x} | \mathcal{F}^e(\mathbf{c}, d)) + \lambda \mathcal{D}(\mathcal{R}(\mathbf{x}, d) || \mathcal{F}^e(\mathbf{c}, d)) \quad (1)$$

where \mathcal{F}^e and \mathcal{F}^d are respectively the encoding and decoding components of a hierarchical generative model; \mathcal{R} is the shared recurrent utterance encoder (the *recognition model*); and \mathcal{D} is a distance function ($L2$ norm).

In turn, with domain description batches, the model maximizes the probability of generating the utterance given its slot-value annotation, both represented as sequences of tokens:

$$\mathcal{L}_{dd} = -\log p_{\mathcal{F}^d}(\mathbf{x} | \mathcal{R}(\mathbf{a}, d)) + \lambda \mathcal{D}(\mathcal{R}(\mathbf{x}, d) || \mathcal{R}(\mathbf{a}, d)) \quad (2)$$

In this multi-task setup, the latent space of \mathcal{R} is shared between both utterances and domain descriptions across all the domains. Moreover, the distance-based loss terms make sure that (a) utterances with similar annotations are closer together in the latent space (Eq. 2), and (b) utterances are closer to their dialogue contexts (Eq. 1) so that their encodings capture some of the contexts’

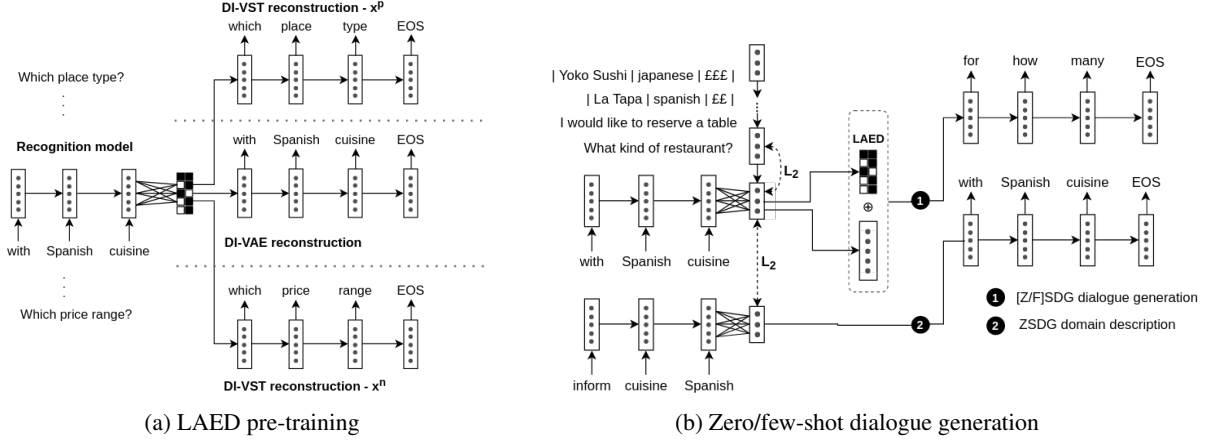


Figure 1: Model architecture. At the pre-training stage (1a), we train the discretized LAED dialogue representation on the Transfer dataset. We then train a zero/few-shot dialogue generation model on SMD with this representation incorporated (1b).

meaning. These properties of the model make it possible to achieve better cross-domain generalization.

3.2 Unsupervised representation learning

As was the case with ZSDG, robust representation learning helps achieve better generalization across domains. The most widely-adopted way to train better representations has been to leverage a greater data source. In this work, we consider unsupervised, variational autoencoder-based (VAE) representation learning on a large dataset of unannotated dialogues. The specific approach we refer to is the Latent Action Encoder-Decoder (LAED) model of (Zhao et al., 2018). LAED is a variant of VAE with two modifications: (1) an optimization objective augmented with mutual information between the input and the latent variable for better and more stable learning performance, and (2) discretized latent variable for the interpretability of the resulting latent actions. Just as in ZSDG, LAED is a hierarchical encoder-decoder model with the key component being a discrete-information (DI) utterance-level VAE. Two versions of this model are introduced, with respective optimization objectives:

$$\mathcal{L}_{DI-VAE} = \mathbb{E}_{q_{\mathcal{R}}(z|x)p(x)} [\log p_{\mathcal{G}}(\mathbf{x} | z)] - KL(q(z)||p(z)) \quad (3)$$

$$\mathcal{L}_{DI-VST} = \mathbb{E}_{q_{\mathcal{R}}(z|x)p(x)} [\log p_{\mathcal{G}}^n(\mathbf{x}_n | z)p_{\mathcal{G}}^p(\mathbf{x}_p | z)] - KL(q(z)||p(z)) \quad (4)$$

where \mathcal{R} and \mathcal{G} are recognition and generation components respectively, \mathbf{x} is the model’s input, z is the latent variable, and $p(z)$ and $q(z)$ are respectively prior and posterior distributions of z .

DI-VAE works in a standard VAE fashion reconstructing the input \mathbf{x} itself, while DI-VST follows the idea of Variational Skip-Thought reconstructing the input’s previous and next contexts: $\{\mathbf{x}_n, \mathbf{x}_p\}$. As reported by the authors, the two models capture different aspects of utterances, i.e. DI-VAE reconstructs specific words within an utterance, whereas DI-VST captures the overall intent better — see the visualization in Figure 1a.

3.3 Proposed models¹

In our approach, we simplify the ZSDG setup by not using any explicit domain descriptions, therefore we only work with ‘dialogue’ batches. We also make use of Knowledge Base information without loss of generality (see Section 5) — thus we work with data of the form $\{c, \mathbf{x}, \mathbf{k}, d\}$ where \mathbf{k} is the KB information. We refer to this model as Few-Shot Dialogue Generation, or *FSDG*.

For learning a reusable dialogue representation, we use an external multi-domain dialogue dataset, the Transfer dataset (see Section 4).

We perform a preliminary training stage on it where we train 2 LAED models, both DI-VAE and DI-VST. Then, at the main training stage, we use the hierarchical encoders of both models and incorporate them with FSDG’s decoder. Thus, we have the following encoding function (which is

¹Code is available at https://bit.ly/fsdg_sigdial2019

Model \ Domain	Navigation		Weather		Schedule	
	BLEU, %	Entity F1, %	BLEU, %	Entity F1, %	BLEU, %	Entity F1, %
ZSDG	5.9	14.0	8.1	31	7.9	36.9
NLU_ZSDG	6.1 ± 2.2	12.7 ± 3.3	5.0 ± 1.6	16.8 ± 6.7	6.0 ± 1.7	26.5 ± 5.4
NLU_ZSDG+LAED	7.9 ± 1	12.3 ± 2.9	8.7 ± 0.6	21.5 ± 6.2	8.3 ± 1	20.7 ± 4.8
FSDG@1%	6.0 ± 1.8	9.8 ± 4.8	6.9 ± 1.1	22.2 ± 10.7	5.5 ± 0.8	25.6 ± 8.2
FSDG@3%	7.9 ± 0.7	11.8 ± 4.4	9.6 ± 1.8	39.8 ± 7	8.2 ± 1.1	34.8 ± 4.4
FSDG@5%	8.3 ± 1.3	15.3 ± 6.3	11.5 ± 1.6	38.0 ± 10.5	9.7 ± 1.4	37.6 ± 8.0
FSDG@10%	9.8 ± 0.8	19.2 ± 3.2	12.9 ± 2.4	40.4 ± 11.0	12.0 ± 1.0	38.2 ± 4.2
FSDG+VAE@1%	3.6 ± 2.6	9.3 ± 4.1	6.8 ± 1.3	23.2 ± 10.1	4.6 ± 1.6	28.9 ± 7.3
FSDG+VAE@3%	6.9 ± 1.9	15.6 ± 5.8	9.5 ± 2.6	32.2 ± 11.8	6.6 ± 1.7	34.8 ± 7.7
FSDG+VAE@5%	7.8 ± 1.9	12.7 ± 4.2	10.1 ± 2.1	40.3 ± 10.4	8.2 ± 1.7	34.2 ± 8.7
FSDG+VAE@10%	9.0 ± 2.0	18.0 ± 5.8	12.9 ± 2.2	40.1 ± 7.6	11.6 ± 1.5	39.9 ± 6.9
FSDG+LAED@1%	7.1 ± 0.8*	10.1 ± 4.5	10.6 ± 2.1*	31.4 ± 8.1*	7.4 ± 1.2	29.1 ± 6.6
FSDG+LAED@3%	9.2 ± 0.8	14.5 ± 4.8*	13.1 ± 1.7	40.8 ± 6.1	9.2 ± 1.2*	32.7 ± 6.1
FSDG+LAED@5%	10.3 ± 1.2	15.6 ± 4.5	14.5 ± 2.2	40.9 ± 8.6	11.8 ± 1.9	37.6 ± 6.1*
FSDG+LAED@10%	12.3 ± 0.9	17.3 ± 4.5	17.6 ± 1.9	47.5 ± 6.0	15.2 ± 1.6	38.7 ± 8.4

Table 1: Evaluation results. Marked with asterisks are individual results higher than the ZSDG baseline which are achieved with the minimum amount of training data, and in bold is the model consistently outperforming ZSDG in all domains and metrics with minimum data.

then plugged in to the Eq. 1):

$$\begin{aligned} \mathcal{F}^e(\mathbf{c}, \mathbf{k}, d) = & \mathcal{F}_{DI-VAE}^e(\mathbf{c}, \mathbf{k}, d) \\ & \oplus \mathcal{F}_{DI-VST}^e(\mathbf{c}, \mathbf{k}, d) \\ & \oplus \mathcal{F}_{FSDG}^e(\mathbf{c}, \mathbf{k}, d) \end{aligned} \quad (5)$$

where \oplus is the concatenation operator. We refer to this model as **FSDG+LAED**.

We compare this LAED-augmented model to a similar one, with latent representation trained on the same data but using a regular VAE objective and thus providing regular continuous embeddings (we refer to it as **FSDG+VAE**).

$$\begin{aligned} \mathcal{L}_{VAE} = & \mathbb{E}_{q_{\mathcal{R}}(z|\mathbf{x})}[\log p_{\mathcal{G}}(\mathbf{x} | z)] \\ & - KL(q_{\mathcal{R}}(z)||p(z)) \end{aligned} \quad (6)$$

Finally, in order to explore the original ZSDG setup as much as possible, we also consider its version with automatic Natural Language Understanding (NLU) markup instead of human annotations as domain descriptions. Our NLU annotations include Named Entity Recognizer (Finkel et al., 2005), a date/time extraction library (Chang and Manning, 2012), and a Wikidata entity linker (Pappu et al., 2017). We have models with (NLU_ZSDG+LAED) and without LAED representation (NLU_ZSDG). Our entire setup is shown in Figure 1.

4 Datasets

We use the Stanford Multi-Domain (SMD) human-human goal-oriented dialogue dataset

(Eric et al., 2017) in 3 domains: appointment scheduling, city navigation, and weather information. Each dialogue comes with knowledge base snippet from the underlying domain-specific API.

For LAED training, we use MetaLWOz (Lee et al., 2019), a human-human goal-oriented dialogue corpus specifically designed for various meta-learning and pre-training purposes. It contains conversations in 51 domains with several tasks in each of those. The dialogues are collected using the Wizard-of-Oz method where human participants were given a problem domain and a specific task. No domain-specific APIs or knowledge bases were available for the participants, and in the actual dialogues they were free to use fictional names and entities in a consistent way. The dataset totals more than 40,000 dialogues, with the average length of 11.9 turns.

5 Experimental setup and evaluation

Our few-shot setup is as follows. Given the target domain, we first train LAED models (a dialogue-level DI-VST and an utterance-level DI-VAE, both of the size 10×5) on the MetaLWOz dataset — here we exclude from training every domain that might overlap with the target one.

Next, using the LAED encoders, we train a Few-Shot Dialogue Generation model on all the SMD source domains. We use a random sample (1% to 10%) of the target domain utterances together with their contexts as seed data.

We incorporate KB information into our model by simply serializing the records and prepending

them to the dialogue context, ending up with a setup similar to CopyNet in (Eric et al., 2017).

For the NLU_ZSDG setup, we use 1000 random seed utterances from each source domain and 200 utterances from the target domain².

For evaluation, we follow the approach of (Zhao and Eskénazi, 2018) and report BLEU and Entity F1 scores — means/variances over 10 runs.

6 Results and discussion

Our results are shown in Table 1. Our objective here is maximum accuracy with minimum training data required, and it can be seen that few-shot models with LAED representation are the best performing models for this objective. While the improvements can already be seen with simple FSDG, the use of LAED representation helps to significantly reduce the amount of in-domain training data needed: in most cases, the state-of-the-art results are attained with as little as 3% of in-domain data. At 5%, we see that FSDG+LAED consistently improves upon all other models in every domain, either by increasing the mean accuracy or by decreasing the variation. In SMD, with its average dialogue length of 5.25 turns (see Table 4), 5% of training dialogues amounts to approximately 200 in-domain training utterances. In contrast, the ZSDG setup used approximately 150 *annotated* training utterances for each of the 3 domains, totalling about 450 *annotated* utterances. Although in our few-shot approach we use full in-domain dialogues, we end up having a comparable amount of target-domain training data, with the crucial difference that none of those has to be annotated for our approach. Therefore, the method we introduced attains state-of-the-art in both accuracy and data-efficiency.

The results of the ZSDG_NLU setup demonstrate that single utterance annotations, if not domain-specific and produced by human experts, don't provide as much signal as raw dialogues.

The comparison of the setups with different latent representations also gives us some insight: while the VAE-powered FSDG model improves on the baseline in multiple cases, it lacks generalization potential compared to LAED. The reason for that might be inherently more stable training of LAED due to its modified objective function

²The numbers are selected so that the domain description task is kept secondary.

which in turn results in a more informative, generalizable representation.

Finally, we discuss the evaluation metrics. Since we base this paper on the work of (Zhao and Eskénazi, 2018), we have had to fully conform to the metrics they used to enable direct comparison. However, BLEU as the primary evaluation metric, does not necessarily reflect NLG quality in dialogue settings — see examples in Table 2 of the Appendix (see also Novikova et al. (2017)). This is a general issue in dialogue model evaluation since the variability of possible responses equivalent in meaning is very high in dialogue. In future work, instead of using BLEU, we will put more emphasis on the meaning of utterances, for example by using external dialogue act tagging resources, using quality metrics of language generation – e.g. perplexity – as well as more task-oriented metrics like Entity F1. We expect these to make for more meaningful evaluation criteria.

7 Conclusion and future work

In this paper, we have introduced a technique to achieve state-of-the-art dialogue generation performance in a few-shot setup, without using any annotated data. By leveraging larger, more highly represented dialogue sources and learning robust latent dialogue representations from them, we obtained a model with superior generalization to an underrepresented domain. Specifically, we showed that our few-shot approach achieves state-of-the-art results on the Stanford Multi-Domain dataset while being more data-efficient than the previous best model, by not requiring any data annotation.

Although being state-of-the-art, the accuracy scores themselves still suggest that our technique is not ready for immediate adoption for real-world production purposes, and the task of few-shot generalization to a new dialogue domain remains an area of active research. We expect that such initiatives will be fostered by the release of large dialogue corpora such as MetaLWOz.

In our own future work, we will try and find ways to improve the unsupervised representation in order to increase the transfer potential. Adversarial learning can also be beneficial in the setting of limited data. And apart from improving the model itself, it is necessary to consider an alternative criterion to BLEU-score for adequate evaluation of response generation.

References

- Angel X. Chang and Christopher D. Manning. 2012. [Sutime: A library for recognizing and normalizing time expressions](#). In *Proceedings of the Eighth International Conference on Language Resources and Evaluation, LREC 2012, Istanbul, Turkey, May 23-25, 2012*, pages 3735–3740.
- Kyunghyun Cho, Bart van Merriënboer, Çağlar Gülçehre, Dzmitry Bahdanau, Fethi Bougares, Holger Schwenk, and Yoshua Bengio. 2014. [Learning phrase representations using RNN encoder-decoder for statistical machine translation](#). In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing, EMNLP 2014, October 25-29, 2014, Doha, Qatar, A meeting of SIGDAT, a Special Interest Group of the ACL*, pages 1724–1734.
- J. Deng, W. Dong, R. Socher, L.-J. Li, K. Li, and L. Fei-Fei. 2009. ImageNet: A Large-Scale Hierarchical Image Database. In *CVPR09*.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2018. [BERT: pre-training of deep bidirectional transformers for language understanding](#). *CoRR*, abs/1810.04805.
- Mihail Eric, Lakshmi Krishnan, Francois Charette, and Christopher D. Manning. 2017. [Key-value retrieval networks for task-oriented dialogue](#). In *Proceedings of the 18th Annual SIGdial Meeting on Discourse and Dialogue, Saarbrücken, Germany, August 15-17, 2017*, pages 37–49.
- A. Eshghi, M. Purver, and Julian Hough. 2011. Dylan: Parser for dynamic syntax. Technical report, Queen Mary University of London.
- Arash Eshghi. 2015. DS-TTR: An incremental, semantic, contextual parser for dialogue. In *Proceedings of Semdial 2015 (goDial), the 19th workshop on the semantics and pragmatics of dialogue*.
- Arash Eshghi, Igor Shalymov, and Oliver Lemon. 2017. [Bootstrapping incremental dialogue systems from minimal data: the generalisation power of dialogue grammars](#). In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing, EMNLP 2017, Copenhagen, Denmark, September 9-11, 2017*, pages 2220–2230.
- Jenny Rose Finkel, Trond Grenager, and Christopher D. Manning. 2005. [Incorporating non-local information into information extraction systems by gibbs sampling](#). In *ACL 2005, 43rd Annual Meeting of the Association for Computational Linguistics, Proceedings of the Conference, 25-30 June 2005, University of Michigan, USA*, pages 363–370.
- Matthew Henderson, Blaise Thomson, and Jason D. Williams. 2014. [The third dialog state tracking challenge](#). In *2014 IEEE Spoken Language Technology Workshop, SLT 2014, South Lake Tahoe, NV, USA, December 7-10, 2014*, pages 324–329.
- Sepp Hochreiter and Jürgen Schmidhuber. 1997. [Long short-term memory](#). *Neural Computation*, 9(8):1735–1780.
- Dimitrios Kalatzis, Arash Eshghi, and Oliver Lemon. 2016. [Bootstrapping incremental dialogue systems: using linguistic knowledge to learn from minimal data](#). In *Proceedings of the NIPS 2016 workshop on Learning Methods for Dialogue*, Barcelona.
- Diederik P. Kingma and Jimmy Ba. 2014. [Adam: A method for stochastic optimization](#). *CoRR*, abs/1412.6980.
- Sungjin Lee, Hannes Schulz, Adam Atkinson, Jianfeng Gao, Kaheer Suleman, Layla El Asri, Mahmoud Adada, Minlie Huang, Shikhar Sharma, Wendy Tay, and Xiujun Li. 2019. [Multi-domain task-completion dialog challenge](#).
- Tomas Mikolov, Ilya Sutskever, Kai Chen, Gregory S. Corrado, and Jeffrey Dean. 2013. [Distributed representations of words and phrases and their compositionality](#). In *Advances in Neural Information Processing Systems 26: 27th Annual Conference on Neural Information Processing Systems 2013. Proceedings of a meeting held December 5-8, 2013, Lake Tahoe, Nevada, United States.*, pages 3111–3119.
- Jekaterina Novikova, Ondrej Dusek, Amanda Cercas-Curry, and Verena Rieser. 2017. [Why we need new evaluation metrics for nlg](#). In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing (EMNLP)*.
- Aasish Pappu, Roi Blanco, Yashar Mehdad, Amanda Stent, and Kapil Thadani. 2017. [Lightweight multilingual entity extraction and linking](#). In *Proceedings of the Tenth ACM International Conference on Web Search and Data Mining, WSDM 17, New York, NY, USA. ACM*.
- Matthew E. Peters, Mark Neumann, Mohit Iyyer, Matt Gardner, Christopher Clark, Kenton Lee, and Luke Zettlemoyer. 2018. [Deep contextualized word representations](#). In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, NAACL-HLT 2018, New Orleans, Louisiana, USA, June 1-6, 2018, Volume 1 (Long Papers)*, pages 2227–2237.
- Igor Shalymov, Arash Eshghi, and Oliver Lemon. 2017. [Challenging neural dialogue models with natural data: Memory networks fail on incremental phenomena](#). In *Proceedings of the 21st Workshop on the Semantics and Pragmatics of Dialogue, SemDial 2017*.
- Sainbayar Sukhbaatar, Arthur Szlam, Jason Weston, and Rob Fergus. 2015. [End-to-end memory networks](#). In *Advances in Neural Information Processing Systems 28: Annual Conference on Neural Information Processing Systems 2015, December 7-12, 2015, Montreal, Quebec, Canada*, pages 2440–2448.

- Stefan Ultes, Pawel Budzianowski, Iñigo Casanueva, Lina Maria Rojas-Barahona, Bo-Hsiang Tseng, Yen-Chen Wu, Steve J. Young, and Milica Gasic. 2018. [Addressing objects and their relations: The conversational entity dialogue model](#). In *Proceedings of the 19th Annual SIGdial Meeting on Discourse and Dialogue, Melbourne, Australia, July 12-14, 2018*, pages 273–283.
- Tsung-Hsien Wen, Yishu Miao, Phil Blunsom, and Steve J. Young. 2017. [Latent intention dialogue models](#). In *Proceedings of the 34th International Conference on Machine Learning, ICML 2017, Sydney, NSW, Australia, 6-11 August 2017*, pages 3732–3741.
- Tiancheng Zhao and Maxine Eskénazi. 2018. [Zero-shot dialog generation with cross-domain latent actions](#). In *Proceedings of the 19th Annual SIGdial Meeting on Discourse and Dialogue, Melbourne, Australia, July 12-14, 2018*, pages 1–10.
- Tiancheng Zhao, Kyusong Lee, and Maxine Eskénazi. 2018. [Unsupervised discrete sentence representation learning for interpretable neural dialog generation](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics, ACL 2018, Melbourne, Australia, July 15-20, 2018, Volume 1: Long Papers*, pages 1098–1107.

Domain	Context	Gold response	Predicted response
schedule	<usr> Remind me to take my pills	Ok setting your medicine appointment for 7pm	Okay, setting a reminder to take your pills at 7 pm.
	<sys> What time do you need to take your pills?		
	<usr> I need to take my pills at 7 pm.		
navigate	<usr> Find the address to a hospital	Have a good day	No problem.
	<sys> Stanford Express Care is at 214 El Camino Real.		
	<usr> Thank you.		
navigate	<usr> What is the weather forecast for the weekend?	For what city would you like to know that?	For what city would you like the weekend forecast for?

Table 2: Selected FSDG+LAED model’s responses

Domain	#Dialogues	Domain	#Dialogues	Domain	#Dialogues
UPDATE_CALENDAR	1991	GUINNESS_CHECK	1886	ALARM_SET	1681
SCAM_LOOKUP	1658	PLAY_TIMES	1601	GAME_RULES	1590
CONTACT_MANAGER	1483	LIBRARY_REQUEST	1339	INSURANCE	1299
HOME_BOT	1210	HOW_TO_BASIC	1086	CITY_INFO	965
TIME_ZONE	951	TOURISM	935	SHOPPING	903
BUS_SCHEDULE_BOT	898	CHECK_STATUS	784	WHAT_IS_IT	776
STORE_DETAILS	737	APPOINTMENT_REMINDER	668	PRESENT_IDEAS	664
GEOGRAPHY	653	SKI_BOT	607	MOVIE_LISTINGS	607
UPDATE_CONTACT	581	ORDER_PIZZA	577	EDIT_PLAYLIST	574
SPORTS_INFO	561	BOOKING_FLIGHT	555	WEATHER_CHECK	551
EVENT_RESERVE	539	RESTAURANT_PICKER	535	LOOK_UP_INFO	533
AUTO_SORT	514	QUOTE_OF_THE_DAY_BOT	513	WEDDING_PLANNER	510
MAKE_RESTAURANT_RESERVATIONS	510	AGREEMENT_BOT	507	NAME_SUGGESTER	499
APARTMENT_FINDER	499	HOTEL_RESERVE	497	PHONE_PLAN_BOT	496
DECIDER_BOT	487	VACATION_IDEAS	486	PHONE_SETTINGS	473
POLICY_BOT	447	PROMPT_GENERATOR	446	MUSIC_SUGGESTER	445
PET_ADVICE	426	BANK_BOT	367	CATALOGUE_BOT	288

Table 3: MetaLWOz domains

A Appendices

A.1 Training details

We train our models with the Adam optimizer (Kingma and Ba, 2014) with learning rate 0.001. Our hierarchical models’ utterance encoder is an LSTM cell (Hochreiter and Schmidhuber, 1997) of size 256, and the dialogue-level encoder is a GRU (Cho et al., 2014) of size 512.

Domain \ Statistic	Navigation	Weather	Schedule
Dialogues	800	797	8285
Utterances	5248	4314	3170
Mean length	6.56	5.41	3.83

Table 4: SMD trainset statistics

A.2 MetaLWOz dialogue example

Domain: Ski

User Task: You want to know if there are good ski hills an hour’s drive from your current location.

Bot Task: Tell the user that there are no ski hills in their immediate location.

Bot Hello how may I help you?

User Is there any good ski hills an hour’s drive from my current location?

Bot I’m sorry to inform you that there are no ski hills in your immediate location

User Can you help me find the nearest?

Bot Absolutely! It looks like you’re about 3 hours away from Bear Mountain. That seems to be the closest.

User Hmm.. sounds good

Bot Alright! I can help you get your lift tickets now! When will you be going?

User Awesome! please get me a ticket for 10pax

Bot You’ve got it. Anything else I can help you with?

User None. Thanks again!

Bot No problem!