

Constraint based Knowledge Base Distillation in End-to-End Task Oriented Dialogs

Dinesh Raghu ^{* † 1 2}, Atishya Jain ^{* 1}, Mausam ¹ and Sachindra Joshi ²

¹ IIT Delhi, New Delhi, India

² IBM Research, New Delhi, India

{diraghu1, jsachind}@in.ibm.com, {Atishya.Jain.cs516, mausam}@cse.iitd.ac.in

Abstract

End-to-End task-oriented dialogue systems generate responses based on dialog history and an accompanying knowledge base (KB). Inferring those KB entities that are most relevant for an utterance is crucial for response generation. Existing state of the art scales to large KBs by softly filtering over irrelevant KB information. In this paper, we propose a novel filtering technique that consists of (1) a pairwise similarity based filter that identifies relevant information by respecting the n-ary structure in a KB record. and, (2) an auxiliary loss that helps in separating contextually unrelated KB information. We also propose a new metric – multiset entity F1 which fixes a correctness issue in the existing entity F1 metric. Experimental results on three publicly available task-oriented dialog datasets show that our proposed approach outperforms existing state-of-the-art models.

1 Introduction

Task oriented dialog systems interact with users to achieve specific goals such as restaurant reservation or calendar enquiry. To satisfy a user goal, the system is expected to retrieve necessary information from a knowledge base and convey it using natural language. Recently several end-to-end approaches (Bordes and Weston, 2017; Wu et al., 2018; He et al., 2020b; Madotto et al., 2018) have been proposed for learning these dialog systems.

Inferring the most relevant KB entities necessary for generating the response is crucial for achieving task success. To effectively scale to large KBs, existing approaches (Wen et al., 2018; Wu et al., 2018) distill the KB by softly filtering irrelevant KB information based on the dialog history. For

#	Date	Event	Invitee	Time
1	1 st Feb	dinner	Alex	10am
2	7 th Feb	dentist appointment	-	2pm
3	8 th Feb	dinner	Ana	7pm

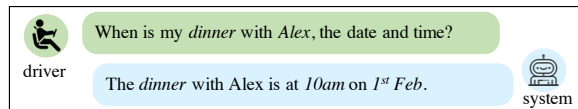


Figure 1: An example dialog between a driver and a system along with the associated knowledge base.

example, in Figure 1 the ideal filtering technique is expected to filter just the row 1 as the driver is requesting information about *dinner* with *Alex*. But existing techniques often filter some irrelevant KB information along with the relevant KB information. For example, in Figure 1 row 3 may also get filtered along with row 1.

Our analysis of the best performing distillation technique (Wu et al., 2018) revealed that embeddings learnt for entities of the same type are quite close to each other. This may be due to entities of the same type often appearing in similar context in history and KB. Such embeddings hurt the overall performance as they reduce the gap between relevant and irrelevant KB records. For example, in Figure 1 row 3 may not get distilled out if *Alex* and *Ana* have similar embeddings.

In this paper, we propose Constraint based knowledge base Distillation NETwork (CDNET), which (1) uses a novel pairwise similarity based distillation computation which distills KB at a record-level, and (2) an auxiliary loss which helps to distill contextually unrelated KB records by enforcing constraints on embeddings of entities of the same type. We noticed the popular entity F1 evaluation metric has a correctness issue when the response contains multi instances of the same entity value. To fix this issue, we propose a new metric called multiset entity F1. We empirically show that CDNET performs either significantly better than or comparable to existing approaches on three pub-

^{*}D. Raghu and A. Jain contributed equally to this work.

[†]D. Raghu is an employee at IBM Research. This work was carried out as part of PhD research at IIT Delhi.

licly available task oriented dialog datasets.

2 Related Work

We first discuss approaches that are closely related to our work. Wu et al. (2018) perform KB distillation but fails to capture the relationship across attributes in KB records. It represents a KB record with multiple attributes as a set of triples (*subject, predicate, object*). This breaks direct connection between record attributes and requires the system to reason over longer inference chains. In Figure 1, if *event* field is used as the key to break the record into triples, then the distillation has to infer that (*dinner, invitee, Alex*), (*dinner, date, 1st Feb*) and (*dinner, time, 10am*) are connected. In contrast, CDNET performs KB distillation by maintaining the attribute relationships. Wen et al. (2018) perform distillation using the similarity between dialog history representation and each attribute representation in a KB record, whereas CDNET uses word based pairwise similarity for distillation.

We now briefly discuss approaches that improve other aspects of task oriented dialogs. He et al. (2020c) and He et al. (2020b) model KBs using Relational GCNs (Schlichtkrull et al., 2018). Raghu et al. (2019) provide support for entities unseen during train. Reddy et al. (2019) improve the ability to reason over KB by respecting the relationships between connected attributes. Qin et al. (2019) restricts the response to contain entities from a single KB record. (Qin et al., 2020) handle multiple domains using shared-private networks and He et al. (2020a) optimize their network on both F1 and BLEU. We are the first to propose a pairwise similarity score for KB distillation and a embedding constraint loss to distill irrelevant KB records.

3 CDNET

CDNET³ has an encoder-decoder architecture that takes as input (1) the dialog history \mathcal{H} , modelled as a sequence of utterances $\{u_i\}_{i=1}^k$ and each utterance u_i as sequence of words $\{w_i^j\}$, and (2) a knowledge base \mathcal{K} with M records $\{r_m\}_{m=1}^M$ and each record r_m has N key-value attribute pairs $\{(k^n, v_m^n)\}_{n=1}^N$. The network generates the system response $Y = (y_1, \dots, y_T)$ one word at a time.

3.1 CDNET Encoder

Context Encoder: The dialog history \mathcal{H} is encoded using a hierarchical encoder (Sordoni et al.,

2015). Each utterance representation \mathbf{u}_i is computed using a Bi-GRU (Schuster and Paliwal, 1997). We denote the hidden state of the j^{th} word in i^{th} utterance as \mathbf{w}_i^j . The context representation \mathbf{c} is generated by passing \mathbf{u}_i s through a GRU.

KB Encoder: We encode the KB using the multi-level memory proposed by Reddy et al. (2019) as its structure allows us to perform distillations over KB records. The KB memory contains two-levels. The first level is a set of KB records. Each KB record is represented as sum of its attributes $\mathbf{r}_m = \sum_n \Phi^e(v_m^n)$, where Φ^e is the embedding matrix. In the second level, each record is represented as a set of attributes. Each attribute is a key-value pair, where the key \mathbf{k}^n is the attribute type embedding and the value \mathbf{v}_m^n is the attribute embedding.

3.2 KB Distillation

The KB distillation module softly filters irrelevant KB records based on the dialog history by computing a distillation distribution (P_d) over the KB records. To compute $P_d = [d_1, \dots, d_M]$, we first score each KB record r_m based on the dialog history \mathcal{H} as follows:

$$s_m = \sum_{w \in \mathcal{H}} \sum_{v_m^n \in r_m} \text{CosSim}(\Phi^e(w), \Phi^e(v_m^n)) \quad (1)$$

where CosSim is the cosine similarity between two vectors. The distillation likelihood d_m for each record r_m then is given by $d_m = \exp(s_m) / \sum_{q=1}^M \exp(s_q)$.

Defining distillation distribution over the KB records rather than KB triples has two main advantages: (1) attributes (such as invitee, event, time and date in Figure 1) in a KB record are directly connected and thus easy to distill, (2) it helps to distill the right records even when the record keys are not unique. In Figure 1, row 3 would be distilled even though it shares the same event name.

3.3 CDNET Decoder

Following Wu et al. (2018), we first generate a sketch response which uses entity type (or sketch) tag in place of an entity. For example, *The @meeting with @invitee is at @time* is generated instead of *The dinner with Alex is at 10pm*. When an entity tag is generated, we choose an entity suggested by the context and KB memory pointers.

Sketch RNN: We use a GRU to generate the sketch response. At each time t , a generate distribution

³<https://github.com/dair-iitd/CDNet>

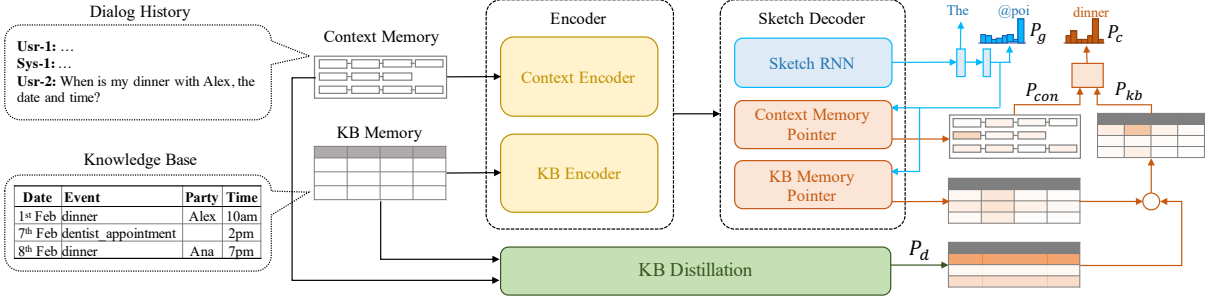


Figure 2: Architecture of CDNET model.

P_g is computed using the decoder hidden state \mathbf{h}_t and an attended summary of the dialog context \mathbf{g}_t . The summary $\mathbf{g}_t = \sum_i \sum_j a_{ij} \mathbf{w}_i^j$, where a_{ij} is the Luong attention (Luong et al., 2015) weights over the context word representations (\mathbf{w}_i^j).

Context Memory Pointer: At each time t , generate the copy distribution over the context P_{con} by performing a multi-hop Luong attention over the context memory. The initial query \mathbf{q}_t^0 is set to \mathbf{h}_t . \mathbf{q}_t^0 is then attended over the context to generate an attention distribution \mathbf{a}^1 and a summarized context \mathbf{g}_t^1 . We represent this as $\mathbf{g}_t^1 = \text{Hop}(\mathbf{q}_t^0, \mathbf{x})$. In the next hop the same process is repeated by updating the query $\mathbf{q}_t^1 = \mathbf{q}_t^0 + \mathbf{g}_t^1$. The attention weights after H hops is used for computing the context pointer P_{con} as follows:

$$P_{con}(y_t = w) = \sum_{ij:w_i^j=w} a_{ij}^H \quad (2)$$

KB Memory Pointer: At each time t , we generate the copy distribution over the KB P_{kb} using (1) Luong attention weight β_m^t over the KB record r_m and (2) Luong attention weight γ_n^t over attribute keys in a record \mathbf{k}^n and (3) the distillation weight d_m over the KB record r_m . The KB pointer P_{kb} is computed as follows:

$$P_{kb}(y_t = w) = \frac{\sum_{mn:v_m^n=w} d_m \beta_m^t \gamma_n^t}{\sum_{mn} d_m \beta_m^t \gamma_n^t} \quad (3)$$

The two copy pointers are combined using a soft gate α (See et al., 2017) to get the final copy distribution P_c as follows,

$$P_c(y_t) = \alpha P_{kb}(y_t) + (1 - \alpha) P_{con}(y_t) \quad (4)$$

3.4 Loss

We guide the distillation module using two auxiliary loss terms: entity constraint loss \mathcal{L}_{ec} and distillation loss \mathcal{L}_d . Often entities of the same type

(e.g., *Ana* and *Alex*) have embeddings similar to each other. As a result, records with similar but unrelated entities are incorrectly assigned a high distillation likelihood. To alleviate this problem, we make the cosine similarity between two entities of the same type to be as low as possible. This is captured by the constraint loss \mathcal{L}_{ec} given by,

$$\mathcal{L}_{ec} = \sum_{(e_a, e_b) \in \mathcal{E}} \text{CosSim}(\Phi^e(e_a), \Phi^e(e_b)) \quad (5)$$

where \mathcal{E} is a set of entity pairs in the KB that belong to the same entity type.

The distillation likelihood d_m of a KB record r_m depends on the similarity between entities in the record and the words mentioned in the dialog context. We compute the distillation loss by defining reference distillation distribution d_m^* as $s_m^* / \sum_{q=1}^M s_q^*$, where s_m^* is the number of times any attribute in r_m occurs in \mathcal{H} and in the gold response. The distillation loss is given by,

$$\mathcal{L}_d = - \sum_{m=1}^M d_m^* \log(d_m) \quad (6)$$

The overall loss function $\mathcal{L} = \mathcal{L}_g + \mathcal{L}_c + \mathcal{L}_{ec} + \mathcal{L}_d$, where \mathcal{L}_g and \mathcal{L}_c are the cross entropy loss on P_g and P_c respectively. Detailed equations are described in Appendix B.

4 Experimental Setup

Datasets: We evaluate our model on three datasets – CamRest (Wen et al., 2017), Multi-WOZ 2.1 (WOZ) (Budzianowski et al., 2018) and Stanford Multi-Domain (SMD) Dataset (Eric et al., 2017).

Baselines: We compare CDNET against the following baselines: MLM (Reddy et al., 2019), DSR (Wen et al., 2018), GLMP (Wu et al., 2018), Entity Consistent (Qin et al., 2019), EER (He et al., 2020c), FG2Seq (He et al., 2020b), TTOS (He et al., 2020a) and DFNet (Qin et al., 2020).

Training Details: CDNET is trained end to end using Adam optimizer (Kingma and Ba, 2014). The

	Utterance	Set	MultiSet
Gold	for which one? I have two, one on the <i>8th</i> at <i>11am</i> and one on <i>wednesday</i> at <i>11am</i>	{ <i>8th, 11am, wednesday</i> }	{ <i>8th, 11am, wednesday, 11am</i> }
Pred-1	your appointment is on <i>8th</i> at <i>11am</i>	{ <i>8th, 11am</i> }	{ <i>8th, 11am</i> }
Pred-2	your appointment is on <i>8th</i> on <i>8th</i> on <i>8th</i> and on <i>8th</i>	{ <i>8th</i> }	{ <i>8th 8th, 8th, 8th</i> }

Table 1: An example to demonstrate the correctness issue with the Entity F1 metric.

Model	CamRest			SMD			WOZ 2.1		
	BLEU	Ent. F1	MSE F1	BLEU	Ent. F1	MSE F1	BLEU	Ent. F1	MSE F1
DSR (Wen et al., 2018)	18.3	53.6	-	12.7	51.9	-	9.1	30.0	-
GLMP (Wu et al., 2018)	15.1	58.9	57.5	13.9	59.6	59.6	6.9	32.4	-
MLM (Reddy et al., 2019)	15.5	62.1	-	17	54.6	-	-	-	-
Ent. Const. (Qin et al., 2019)	18.5	58.6	-	13.9	53.7	-	-	-	-
TTOS (He et al., 2020a)	20.5	61.5	-	17.4	55.4	-	-	-	-
DFNet (Qin et al., 2020)	-	-	-	14.4	62.7	56.7	9.4	35.1	34.8
EER (He et al., 2020c)	19.2	65.7	65.5	17.2	59.0	55.1	13.6	35.6	35.0
FG2Seq (He et al., 2020b)	20.2	66.4	65.4	16.8	61.1	59.1	14.6	36.5	36.0
CDNET	21.8	68.6	68.4	17.8	62.9	62.9	11.9	38.7	38.6

Table 2: Performance of CDNET and baselines on the CamRest, SMD and Multi-WOZ 2.1 datasets.

embedding dimensions of the hidden states of encoder and decoder GRU are set to 200 and 100 respectively. Word embeddings are initialized with pre-trained 200d GloVe embeddings (Pennington et al., 2014). Words not in GloVe are initialized using Glorot uniform distribution (Glorot and Bengio, 2010). The dropout rate is set to 0.2 and teacher forcing ratio set to 0.9. The best hyper-parameter setting for each dataset and other training details are reported in the Appendix A.

Evaluation Metrics: We measure the performance of all the models using BLEU (Papineni et al., 2002), our proposed multiset entity F1 and for completeness the previously used entity F1 (Wu et al., 2018).

MultiSet Entity F1 (MSE F1): The entity F1 is used to measure the model’s ability to predict relevant entities from the KB. It is computed by micro averaging over the set of entities in the gold responses and the set of entities in the predicted responses. This metric suffers from two main problems. First, when the gold response has multiple instances of the same entity value, it is accounted for just once in the set representation. For example, in Table 1 the entity value *11am* occurs twice in the gold response but accounted for just once in the *set* representation. As a result the recall computation does not penalize the prediction *pred-1* for missing an instance of *11am*. Second, the existing metric fails to penalize models that stutter. For example, in Table 1 the precision of *pred-2*

is not penalized for repeating the entity value *8th*. We propose a simple modification to the entity F1 metric to fix these correctness issues. The modified metric, named MultiSet Entity F1, is computed by micro averaging over the *multiset* of entities rather than a *set*. As multisets allow multiple instances of same entity values, it (1) accounts for the same entity value mentioned more than once in the gold by penalizing recall for missing any instances and (2) accounts for models that stutters by penalizing the precision.

5 Results

The results are shown in Table 2. On CamRest and SMD, CDNET outperforms existing models in both MSE F1 and BLEU. On WOZ, CDNET achieves best only in MSE F1. We observed that the responses generated by CDNET on WOZ were appropriate, but did not have good lexical overlap with the gold responses. To investigate this further, we perform a human evaluation of the responses predicted by CDNET, FG2Seq and EER.

Human Evaluation: We conduct a human evaluation to assess two dimensions of generated responses: (1) *Appropriateness*: how useful are the responses for the given dialog context and KB, and (2) *Naturalness*: how human-like are the predicted responses. We randomly sampled 75 dialogs from each of the three datasets and requested two judges to evaluate on a Likert scale (Likert, 1932). The results are summarized in Table 3. CDNET outper-

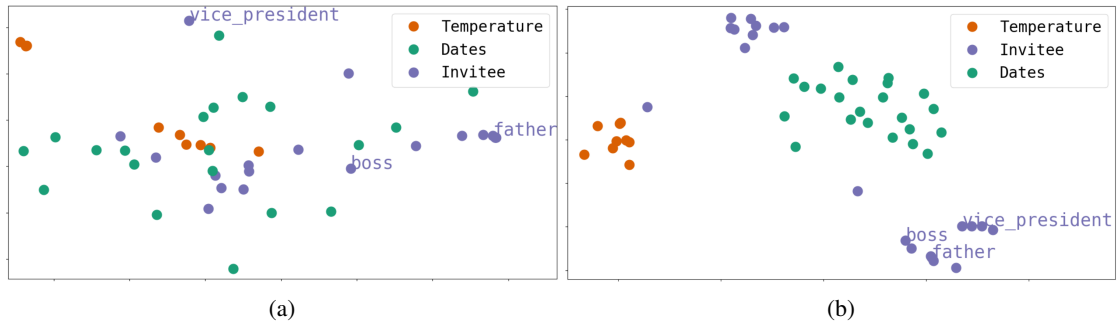


Figure 3: T-sne plots of entity embeddings from SMD of (a) CDNET & (b) GLMP.

Model	Appropriateness			Naturalness		
	SMD	Cam	WoZ	SMD	Cam	WoZ
EER	2.9	3.8	3.4	3.6	4.2	4.0
FG2Seq	3.1	3.7	3.7	3.9	4.3	4.0
CDNET	3.6	4.1	3.9	3.7	4.3	4.1

Table 3: Human Evaluation of CDNET on the CamRest, SMD and Multi-WOZ 2.1 datasets.

Model	CamRest		SMD	
	BLEU	MSE F1	BLEU	MSE F1
CDNET	21.8	68.4	17.8	62.9
No \mathcal{L}_{ec}	19.2	65.4	17.4	62.2
Naive Dist.	15.0	64.2	16.9	60.6
Entry-Level Attn.	16.2	62.0	17.1	59.4

Table 4: Ablation study of CDNET on the CamRest and SMD datasets.

forms both FG2Seq and EER on *appropriateness* across all three datasets. Despite having a lower BLEU score on WOZ, CDNET performs in-par with the other two baselines on *naturalness*.

Ablation Study: We perform an ablation study by defining three variants. Table 4 shows the MSE F1 and BLEU for the two settings on CamRest and SMD datasets. (1) We remove the entity constraint loss \mathcal{L}_{ec} from the overall loss \mathcal{L} . (2) We replace our pairwise similarity based score s_m used for KB distillation with the global pointer score ($s_m = \mathbf{r}_m^T \cdot \mathbf{c}$) proposed by (Wu et al., 2018). We refer to this setting as naive distillation. (3) We replace our pairwise similarity based score s_m with the entry-level attention proposed by (Wen et al., 2018). We see that both our contributions: pairwise similarity scorer for computing distillation distribution and the entity constraint loss contribute to the overall performance.

Discussion: We now discuss the effect of the entity constraint loss \mathcal{L}_{ec} on the KB entity embeddings. Figure 3 shows the t-SNE plot (Van der Maaten and Hinton, 2008) of entity embeddings of CD-

NET and GLMP where entities of the same type are represented using the same colour. We see that entities of the same type (e.g. *father* and *boss* of the type *invitees*) are clustered together in embedding space of GLMP, while they are distributed across the space in CDNET. This shows that the entity constraint loss has helped reduce the embedding similarity between entities of the same type and ensures KB records with similar but unrelated entities are filtered by the KB distillation. Visualization of distillation distribution helping identify relevant KB entities is shown in Appendix C.

6 Conclusion

We propose CDNET for learning end-to-end task oriented dialog system. CDNET performs KB distillation at the level of KB records, thereby respecting the relationships between the connected attributes. CDNET uses a pairwise similarity based score function to better distill the relevant KB records. By defining constraints over embeddings of entities of the same type, CDNET filters out contextually unrelated KB records. We propose a simple modification to the entity F1 metric that helps fix correctness issues. We refer to the new metric as multiset entity F1. CDNET significantly outperforms existing approaches on multiset entity F1 and *appropriateness*, while being comparable on *naturalness* and BLEU. We release the code for further research.

Acknowledgments

This work is supported by IBM AI Horizons Network grant, an IBM SUR award, grants by Google, Bloomberg and IMG, a Visvesvaraya faculty award by Govt. of India, and the Jai Gupta chair fellowship by IIT Delhi. We thank the IIT Delhi HPC facility for computational resources.

References

- Antoine Bordes and Jason Weston. 2017. Learning end-to-end goal-oriented dialog. In *International Conference on Learning Representations*.
- Paweł Budzianowski, Tsung-Hsien Wen, Bo-Hsiang Tseng, Iñigo Casanueva, Stefan Ultes, Osman Ramadan, and Milica Gasic. 2018. Multiwoz-a large-scale multi-domain wizard-of-oz dataset for task-oriented dialogue modelling. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 5016–5026.
- Mihail Eric, Lakshmi Krishnan, Francois Charette, and Christopher D. Manning. 2017. Key-value retrieval networks for task-oriented dialogue. In *Dialog System Technology Challenges, Saarbrücken, Germany, August 15-17, 2017*, pages 37–49.
- Xavier Glorot and Yoshua Bengio. 2010. Understanding the difficulty of training deep feedforward neural networks. In *Proceedings of the thirteenth international conference on artificial intelligence and statistics*, pages 249–256. JMLR Workshop and Conference Proceedings.
- Wanwei He, Min Yang, Rui Yan, Chengming Li, Ying Shen, and Ruifeng Xu. 2020a. [Amalgamating knowledge from two teachers for task-oriented dialogue system with adversarial training](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 3498–3507, Online. Association for Computational Linguistics.
- Zhenhao He, Yuhong He, Qingyao Wu, and Jian Chen. 2020b. [Fg2seq: Effectively encoding knowledge for end-to-end task-oriented dialog](#). In *ICASSP 2020-2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 8029–8033. IEEE.
- Zhenhao He, Jiachun Wang, and Jian Chen. 2020c. Task-oriented dialog generation with enhanced entity representation. *Proc. Interspeech 2020*, pages 3905–3909.
- Diederik P Kingma and Jimmy Lei Ba. 2014. Adam: A method for stochastic optimization. In *Proc. 3rd Int. Conf. Learn. Representations*.
- Rensis Likert. 1932. A technique for the measurement of attitudes. *Archives of psychology*.
- Thang Luong, Hieu Pham, and Christopher D. Manning. 2015. Effective approaches to attention-based neural machine translation. In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*, pages 1412–1421. Association for Computational Linguistics.
- Laurens Van der Maaten and Geoffrey Hinton. 2008. Visualizing data using t-sne. *Journal of machine learning research*, 9(11).
- A. Madotto, CS. Wu, and P. Fung. 2018. Mem2seq: Effectively incorporating knowledge bases into end-to-end task-oriented dialog systems. In *Proceedings of 56th Annual Meeting of the Association for Computational Linguistics*. Association for Computational Linguistics.
- 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 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.
- Libo Qin, Yijia Liu, Wanxiang Che, Haoyang Wen, Yangming Li, and Ting Liu. 2019. Entity-consistent end-to-end task-oriented dialogue system with kb retriever. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 133–142.
- Libo Qin, Xiao Xu, Wanxiang Che, Yue Zhang, and Ting Liu. 2020. [Dynamic fusion network for multi-domain end-to-end task-oriented dialog](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 6344–6354, Online. Association for Computational Linguistics.
- Dinesh Raghu, Nikhil Gupta, and Mausam. 2019. [Disentangling Language and Knowledge in Task-Oriented Dialogs](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 1239–1255, Minneapolis, Minnesota. Association for Computational Linguistics.
- Revanth Gangi Reddy, Danish Contractor, Dinesh Raghu, and Sachindra Joshi. 2019. Multi-level memory for task oriented dialogs. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 3744–3754.
- Michael Sejr Schlichtkrull, Thomas N Kipf, Peter Bloem, Rianne van den Berg, Ivan Titov, and Max Welling. 2018. Modeling relational data with graph convolutional networks. In *ESWC*.
- Mike Schuster and Kuldip K Paliwal. 1997. Bidirectional recurrent neural networks. *IEEE transactions on Signal Processing*, 45(11):2673–2681.
- Abigail See, Peter J Liu, and Christopher D Manning. 2017. Get to the point: Summarization with pointer-generator networks. In *Proceedings of the 55th Annual Meeting of the Association for Computational*

Linguistics (Volume 1: Long Papers), pages 1073–1083.

Alessandro Sordani, Yoshua Bengio, Hossein Vahabi, Christina Lioma, Jakob Grue Simonsen, and Jian-Yun Nie. 2015. A hierarchical recurrent encoder-decoder for generative context-aware query suggestion. In *Proceedings of the 24th ACM International on Conference on Information and Knowledge Management*, pages 553–562.

Haoyang Wen, Yijia Liu, Wanxiang Che, Libo Qin, and Ting Liu. 2018. [Sequence-to-sequence learning for task-oriented dialogue with dialogue state representation](#). In *Proceedings of the 27th International Conference on Computational Linguistics*, pages 3781–3792, Santa Fe, New Mexico, USA. Association for Computational Linguistics.

TH Wen, D Vandyke, N Mrkšić, M Gašić, LM Rojas-Barahona, PH Su, S Ultes, and S Young. 2017. A network-based end-to-end trainable task-oriented dialogue system. In *15th Conference of the European Chapter of the Association for Computational Linguistics, EACL 2017-Proceedings of Conference*, volume 1, pages 438–449.

Chien-Sheng Wu, Richard Socher, and Caiming Xiong. 2018. Global-to-local memory pointer networks for task-oriented dialogue. In *International Conference on Learning Representations*.

A Training Details:

All the hyper parameters are finalised after a grid search over the dev set. We sample learning rates (LR) from $\{2.5 \times 10^{-4}, 5 \times 10^{-4}, 10^{-4}\}$. The Disentangle Label Dropout (DLD) rate (Raghu et al., 2019) is sampled from $\{0.0, 0.05, 0.10, 0.15, 0.20\}$. The number of hops H in the response decoder is sampled from $\{1, 3, 5\}$. We ran each hyperparameter setting 10 times and use the setting with the best validation entity F1. The best performing hyperparameters for all datasets are listed in Table 5.

Dataset	Hops	DLD	LR	Val MSE F1
CamRest	1	0%	0.0005	68.6
SMD	3	5%	0.00025	60.4
WoZ 2.1	3	0%	0.00025	34.3

Table 5: Best performing hyperparameters along with the best validation Entity F1 (Val Ent. F1) achieved for the three datasets.

All experiments were run on a single Nvidia V100 GPU with 32GB of memory. CDNET has an average runtime of 3 hours (6 min per epoch), 10 hours (20 min per epoch) and 24 hours (36 min per epoch) on CamRest, SMD and WOZ respectively. CDNET has a total of 2.8M trainable parameters (400K for embedding matrix, 720K for context encoder, 240k for the sketch RNN and 1440k for the Memory pointers).

B Detailed Equations

In this section, we describe the details of context encoder, CDNETdecoder and the loss.

B.1 Context Encoder

Given a dialog history \mathcal{H} we compute the utterance representation \mathbf{u}_i and context representation \mathbf{c} as follows:

$$\mathbf{u}_i = \text{BiGRU}(\Phi^e(w_i^1), \dots, \Phi^e(w_i^{\tau_i})) \quad (7)$$

$$\mathbf{c} = \text{GRU}(\mathbf{u}_1, \dots, \mathbf{u}_k) \quad (8)$$

where τ_i is the number of words in utterance u_i and w_i^j is the j^{th} word in the i^{th} utterance.

B.2 CDNET Decoder

Let \mathbf{h}_t and y_t be the hidden state and the predicted word at time t respectively. The hidden state is computed as follows,

$$\mathbf{h}_t = \text{GRU}(\Phi^e(y_{t-1}), \mathbf{h}_{t-1}) \quad (9)$$

Now, we compute multi-hop Luong attention over the words representations \mathbf{w}_i^j in the context memory. We set the initial query \mathbf{q}_t^0 to \mathbf{h}_t and then apply Luong attention as follows:

$$a_{ij}^1 = \frac{\exp(W_1 \tanh(W_2[\mathbf{w}_i^j, \mathbf{q}_t^0]))}{\sum_{ij} \exp(W_1 \tanh(W_2[\mathbf{q}_t^0, \mathbf{w}_i^j]))} \quad (10)$$

where W_1, W_2 are trainable parameters. We then compute the summarized context representation \mathbf{g}_t^1 and the next hop query as follows:

$$\mathbf{g}_t^1 = \sum_{ij} a_{ij}^1 \mathbf{w}_i^j \quad (11)$$

$$\mathbf{q}_t^1 = \mathbf{q}_t^0 + \mathbf{g}_t^1 \quad (12)$$

We repeat this for H hops. The attention vector after H hop is represented \mathbf{a}^H . The generate distribution P_g is given by:

$$P_g = \text{Softmax}(W_3[\mathbf{h}_t, \mathbf{g}_t^1] + b_1) \quad (13)$$

where W_3 and b_1 are trainable parameters. The context copy distribution P_{con} is computed as follows:

$$P_{con}(y_t = w) = \sum_{ij: w_i^j = w} a_{ij}^H \quad (14)$$

The KB copy distribution P_{kb} is given by,

$$\beta_m^t = \text{Softmax}(W_4 \tanh(W_5[\mathbf{g}_t^H, \mathbf{h}_t, \mathbf{r}_m])) \quad (15)$$

$$\gamma_n^t = \text{Softmax}(W_6 \tanh(W_7[\mathbf{g}_t^H, \mathbf{h}_t, \mathbf{k}^n])) \quad (16)$$

$$P_{kb}(y_t = w) = \frac{\sum_{mn: v_m^n = w} d_m \beta_m^t \gamma_n^t}{\sum_{mn} d_m \beta_m^t \gamma_n^t} \quad (17)$$

where W_4, W_5, W_6 and W_7 are trainable parameters. Now we compute the gate α to combine P_{con} and P_{kb} to get a final copy distribution P_c as follows:

$$\mathcal{K}_t = \sum_m d_m \beta_m^t \mathbf{r}_m \quad (18)$$

$$\alpha = \text{Sigmoid}(W_8[\mathbf{h}_t, \mathbf{g}_t^H, \mathcal{K}_t]) \quad (19)$$

$$P_c = \alpha P_{kb}(y_t) + (1 - \alpha) P_{con}(y_t) \quad (20)$$

where W_8 is a trainable parameter.

B.3 Loss

We compute the cross entropy loss over the generate P_g and copy P_c distribution as follows:

$$\mathcal{L}_g = - \sum_{t=1}^T \log(P_g(y_t)) \quad (21)$$

$$\mathcal{L}_c = - \sum_{t=1}^T \log(P_c(y_t)) \quad (22)$$

C Distillation Visualisation

We show the visualisation of how the KB distillation distribution helps the decoder rectify the incorrect KB memory pointer inference in Figure 4. Figure 5 shows how the KB distillation distribution helps increase the confidence associated with the correct entity in the KB.

D Datasets

We present statistics of SMD, CamRest and WOZ in Table 6.

	SMD	CamRest	WOZ
Train Dialogs	2425	406	1839
Val Dialogs	302	135	117
Test Dialogs	304	135	141

Table 6: Statistics of the three datasets.

E Domain-Wise Results

Table 7 and Table 8 show the domain wise entity F1 scores of SMD and WOZ datasets respectively. We note that CDNET either has the best or the second-best performance in domain wise scores.

F Qualitative Example

Table 9 shows responses predicted by CDNET, EER and FG2Seq for an example from the WOZ

dataset.

G Human Evaluation

Figure 6 shows a screenshot of the task used for collecting human judgements.

Model	BLEU	F1	MSE F1	Cal	Wea	Nav
MLM	17.0	54.6	-	66.7	56	46.9
DSR	12.7	51.9	-	52.1	50.4	52.0
Ent. Const.	13.9	53.7	-	55.6	52.2	54.5
TTOS	17.4	55.4	-	63.5	64.1	45.9
DFNet	14.4	62.7	-	73.1	57.6	57.9
GLMP	13.9	59.6	59.6	70.2	58.0	54.3
EER	17.2	59.0	55.1	71.8	57.8	52.5
FG2Seq	16.8	61.1	59.1	73.3	57.4	56.1
CDNET	17.8	62.9	62.9	75.4	61.3	56.7

Table 7: Domain wise Entity F1 performance of CDNET and baselines on the SMD dataset.

Model	BLEU	F1	MSE F1	Hot	Att	Res
DSR	9.1	30.0	-	27.1	28.0	33.4
DFNet	9.4	35.1	34.8	30.6	28.1	40.9
GLMP	6.9	32.4	-	28.1	24.4	38.4
EER	13.6	35.6	35.0	35.7	43.0	34.2
FG2Seq	14.6	36.5	36.0	34.4	37.2	38.9
CDNET	11.9	38.7	38.6	36.3	38.9	41.7

Table 8: Domain wise Entity F1 performance of CDNET and baselines on WOZ dataset.

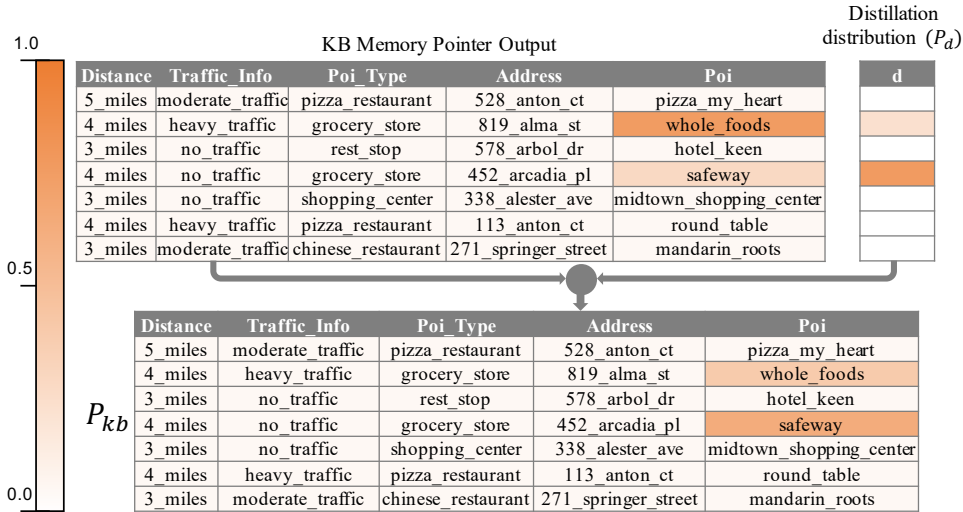
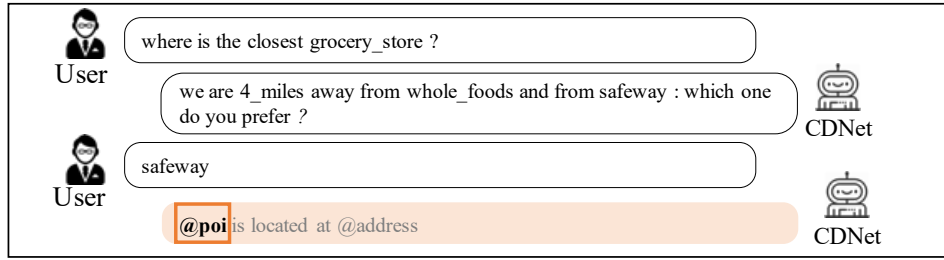


Figure 4: Attention visualization of a decode time step of an example from SMD dataset. P_{kb} corresponds to the sketch tag @poi. P_{kb} is computed by combining the output of the KB memory pointer and the distillation distribution P_d .

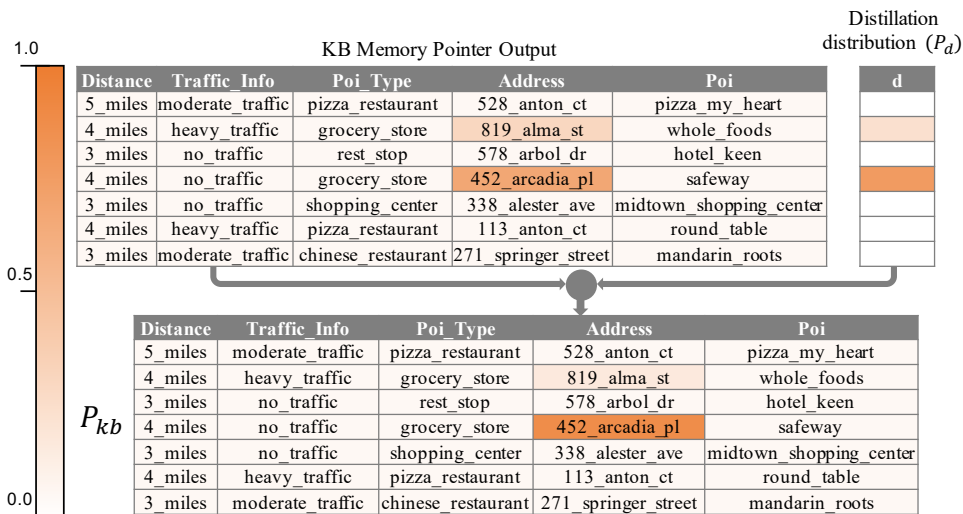
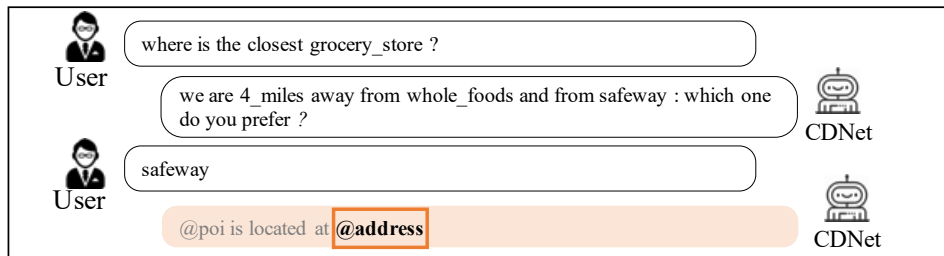


Figure 5: Attention visualization of a decode time step of an example from SMD dataset. P_{kb} corresponds to the sketch tag @address. P_{kb} is computed by combining the output of the KB memory pointer and the distillation distribution P_d .

Address	Area	Choice	Name	Postcode	Pricerange	Ref	Food
regent_street_city_centre	centre	a_lot	pizza_hut_city_centre	cb21ab	cheap	oog9kq5v	italian
4_-6_rose_crescent	centre	a_lot	la_raza	cb23ll	cheap	oog9kq5v	spanish
finders_corner_newmarket_road	east	a_lot	the_missing_sock	cb259aq	cheap	oog9kq5v	international
106_regent_street_city_centre	centre	a_lot	curry_garden	cb21dp	expensive	oog9kq5v	indian

usr-1	i need a cheap spanish restaurant . are there any around ?
sys-1	yes , la_raza is a restaurant serving spanish food in the cheap price range . it is located at 4_-6_rose_crescent . would you like a reservation ?
usr-2	i would like a restaurant in the centre .
sys-2	la_raza is a cheap restaurant serving spanish food in the centre part of town . it ' s address is 4_-6_rose_crescent and it ' s post code is cb2311 .
usr-3	great . can you book me a table for 7 on tuesday at 11:45 and provide me with a reference number ?
Gold	all right , i ' ve made a reservation . your reference number is oog9kq5v . what else can i do for you ?
EER	i ' m sorry , but i was unable to book that for you . would you like to try another restaurant ?
FG2Seq	i ' m sorry , but the la_raza restaurant restaurants . would you like me to find another restaurant ?
CDNET	i have booked you at la_raza . your reference number is oog9kq5v . is there anything else i can help you with ?

Table 9: Example dialog from WOZ along with responses generated by EER, FG2Seq and CDNET.

Address	Distance	Poi	Poi_type	Traffic_info							
580_van_ness_ave	5_miles	toms_house	friends_house	moderate_traffic							
465_arcadia_pl	6_miles	four_seasons	rest_stop	car_collision_nearby							
333_arbol_dr	4_miles	travelers_lodge	rest_stop	car_collision_nearby							
611_ames_ave	4_miles	palo_alto_medical_foundation	hospital	no_traffic							
394_van_ness_ave	6_miles	coupa	coffee_or_tea_place	moderate_traffic							
409_bollard_st	2_miles	willows_market	grocery_store	car_collision_nearby							
91_el_camino_real	2_miles	76	gas_station	road_block_nearby							
5671_barringer_street	6_miles	home	home	heavy_traffic							
User	where is a nearby rest_stop ?										
System	the travelers_lodge is located 4_miles away at 333_arbol_dr would you like directions there ?										
User	yes , please give me directions via the fastest route										
	Appropriateness		Naturalness								
Sample Answer	ok , i am setting gps for 333_arbol_dr										
Model 1	"there is a car_collision_nearby but i sent on your screen another route to reach there "	<input type="radio"/> SD	<input type="radio"/> D	<input type="radio"/> N	<input type="radio"/> A	<input type="radio"/> SA	<input type="radio"/> SD	<input type="radio"/> D	<input type="radio"/> N	<input type="radio"/> A	<input type="radio"/> SA
Model 2	"i sent the info on your screen , drive carefully ! "	<input type="radio"/> SD	<input type="radio"/> D	<input type="radio"/> N	<input type="radio"/> A	<input type="radio"/> SA	<input type="radio"/> SD	<input type="radio"/> D	<input type="radio"/> N	<input type="radio"/> A	<input type="radio"/> SA
Model 3	"setting navigation to 333_arbol_dr now "	<input type="radio"/> SD	<input type="radio"/> D	<input type="radio"/> N	<input type="radio"/> A	<input type="radio"/> SA	<input type="radio"/> SD	<input type="radio"/> D	<input type="radio"/> N	<input type="radio"/> A	<input type="radio"/> SA

Figure 6: A sample human evaluation task used for collecting appropriateness and naturalness of responses generated by three (anonymized) models.