

Zero-Shot Dialogue Disentanglement

by

Self-Supervised Entangled Response Selection

Ta-Chung Chi

Language Technologies Institute
Carnegie Mellon University
tachungc@andrew.cmu.edu

Alexander I. Rudnicky

Language Technologies Institute
Carnegie Mellon University
air@cs.cmu.edu

Abstract

Dialogue disentanglement aims to group utterances in a long and multi-participant dialogue into threads. This is useful for discourse analysis and downstream applications such as dialogue response selection, where it can be the first step to construct a clean context/response set. Unfortunately, labeling all *reply-to* links takes quadratic effort w.r.t the number of utterances: an annotator must check all preceding utterances to identify the one to which the current utterance is a reply. In this paper, we are the first to propose a **zero-shot** dialogue disentanglement solution. Firstly, we train a model on a multi-participant response selection dataset harvested from the web which is not annotated; we then apply the trained model to perform zero-shot dialogue disentanglement. Without any labeled data, our model can achieve a cluster F1 score of 25. We also fine-tune the model using various amounts of labeled data. Experiments show that with only 10% of the data, we achieve nearly the same performance of using the full dataset¹.

1 Introduction

Multi-participant chat platforms such as Messenger and WhatsApp are common on the Internet. While being easy to communicate with others, messages often flood into a single channel, entangling chat history which is poorly organized and difficult to structure. In contrast, Slack provides a thread-opening feature that allows users to manually organize their discussions. It would be ideal if we could design an algorithm to automatically organize an entangled conversation into its constituent threads. This is referred to as the task of *dialogue disentanglement* (Shen et al., 2006; Elsner and Charniak, 2008; Wang and Oard, 2009; Elsner and Charniak, 2011; Jiang et al., 2018; Kummerfeld et al., 2018; Zhu et al., 2020; Li et al., 2020; Yu and Joty, 2020).

¹Code is released at https://github.com/chijames/zero_shot_dialogue_disentanglement

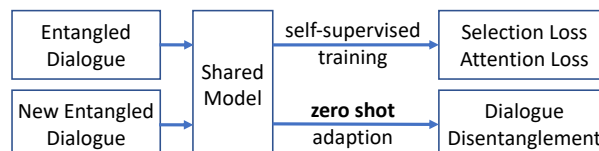


Figure 1: This is the high-level flow of our proposed approach.

Training data for the dialogue disentanglement task is difficult to acquire due to the need for manual annotation. Typically, the data is annotated in the *reply-to* links format, i.e. every utterance is linked to one preceding utterance. The effort is quadratic w.r.t the length of dialogue, partly explaining the sole existence of human-annotated large-scale dataset (Kummerfeld et al., 2018), which was constructed based on the Ubuntu IRC forum. To circumvent the need for expensive labeled data, we aim to train a self-supervised model first then use the model to perform zero-shot dialogue disentanglement. In other words, our goal is to find a task that can learn implicit *reply-to* links without labeled data.

Entangled response selection (Gunasekara et al., 2020) is the task that we will focus on. It is similar to the traditional response selection task, whose goal is to pick the correct next response among candidates, with the difference that its dialogue context consists of multiple topics and participants, leading to a much longer context (avg. 55 utterances). We hypothesize that:

A well-performing model of entangled response selection requires recovery of reply-to links to preceding dialogue.

This is the only way that a model can pick the correct next response given an entangled context. Two challenges are ahead of us:

[com/chijames/zero_shot_dialogue_disentanglement](https://github.com/chijames/zero_shot_dialogue_disentanglement)

- Choosing a design for such a model. Previous work relies on heuristics to filter out utterances to condense context. A model should not rely on heuristics. See §2.3 and 2.5.
- Even though we can train a well-performing model, how should we reveal the links learned implicitly? See §3.4.

Finally, we want to highlight the high practical value of our proposed method. Consider that we have access to a large and unlabeled corpus of chat (e.g. WhatsApp/Messenger) history. The only cost should be training the proposed entangled response selection model with attention supervision using unlabeled data. The trained model is immediately ready for dialogue disentanglement. In summary, the contributions of this work are:

- Show that complex pruning strategies are not necessary for entangled response selection.
- With the proposed objective, the model trained on entangled response selection can perform zero-shot dialogue disentanglement.
- By tuning with 10% of the labeled data, our model achieves comparable performance to that trained using the full dataset.

2 Entangled Response Selection

2.1 Task Description

The dataset we use is DSTC8 subtask-2 (Gunnasekara et al., 2020), which was constructed by crawling the Ubuntu IRC forum. Concretely, given an entangled dialogue context, the model is expected to pick the next response among 100 candidates. The average context length is 55 and the number of speakers is 20 with multiple (possibly relevant) topics discussed concurrently. The context is too long to be encoded by transformer-based models (Devlin et al., 2018; Liu et al., 2019). Despite the existence of models capable of handling long context (Yang et al., 2019; Zaheer et al., 2020; Beltagy et al., 2020), it is difficult to reveal the implicitly learned *reply-to* links as done in §3.4.

2.2 Related Work

To the best of our knowledge, previous works adopt complex heuristics to prune out utterances in the long context (Wu et al., 2020; Wang et al., 2020; Gu et al., 2020; Bertero et al., 2020). For example, keeping the utterances whose speaker is the same as

Model	R@1	R@5	R@10	MRR
Concatenate	51.6	72.3	80.1	61.2
+ aug	64.3	82.9	88.4	72.8
Hierarchical	50.0	72.1	81.4	60.3
+ aug	65.7	84.8	91.8	74.3

Table 1: Test set performance of the entangled response selection task. Concatenate is the model with complex pruning heuristics described in Wu et al. (2020).

or referred to by the candidates. This is problematic for two reasons. 1) The retained context is still noisy as there are multiple speakers present in the candidates. 2) We might accidentally prune out relevant utterances even though they do not share the same speakers. A better solution is to let the model decide which utterances should be retained.

2.3 Model (Solid Arrows in Figure 2)

We use a hierarchical encoder as shown in the middle part of Figure 2. Suppose the input context is $\{U_i\}_{i=1}^n$ and the next response candidate set is $\{C_k\}_{k=1}^m$. For every candidate utterance C_k , we concatenate it with all U_i s. For example, we form n pairs for $k = 1$, $(U_i + C_1)_{i=1}^n$. Then we use BERT as the encoder (φ) to encode pairs and get the last layer embedding of the [CLS] token as \mathbf{V}_i :

$$\mathbf{V}_i = \varphi(U_i + C_1)_{i=1}^n, \quad \forall i \in 1 \dots n \quad (1)$$

$$\mathbf{V}_{n+1} = \varphi(C_k + C_k) \quad (2)$$

While $(C_k + C_k)$ is not necessary for response selection, it is useful later for predicting self-link, which acts as the first utterance of a thread. We will see its role in §3.4. Then we use the output embeddings of a one layer transformer (ψ) with 8 heads to encode contextualized representations:

$$\{\mathbf{V}'_i\}_{i=1}^{n+1} = \psi(\{\mathbf{V}_i\}_{i=1}^{n+1}) \quad (3)$$

To determine relative importance, we use an attention module (\mathcal{A}) to calculate attention scores:

$$v_i = \text{MLP}(\mathbf{V}'_i), \quad \forall i \in 1 \dots n + 1 \quad (4)$$

$$\{\alpha_i\}_{i=1}^{n+1} = \text{softmax}(\{v_i\}_{i=1}^{n+1}) \quad (5)$$

The final predicted score is:

$$s = \text{MLP}\left(\sum_{i=1}^{n+1} \alpha_i \mathbf{V}'_i\right) \quad (6)$$

Note that s should be 1 for C_1 (the correct next response), and otherwise 0 (row 1 of the multi-task loss table in Figure 2). This can be optimized using the binary cross-entropy loss.

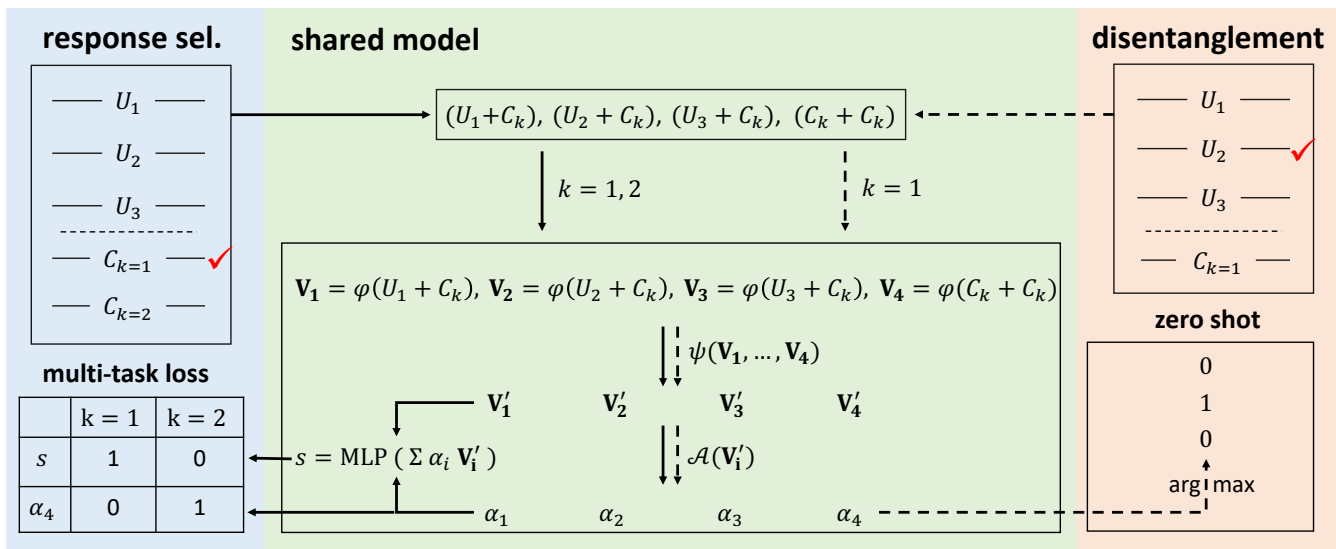


Figure 2: **Solid arrows:** Given an entangled context $U_{1,2,3}$ and $C_{k=1}$ as the correct next response ($C_{k=2}$ is a negative sample), each pair of the concatenated inputs is encoded separately by φ (BERT) to get \mathbf{V}_i . A context-aware model ψ (transformer) is applied over \mathbf{V}_i s to generate contextualized \mathbf{V}'_i . An attention module \mathcal{A} is used to calculate the attention scores α_i and weighted sum s . Model is optimized according to the target values for s and α_4 in the multi-task loss table. **Dashed arrows:** Given another entangled context, we know that the current utterance $C_{k=1}$ is replying to U_2 by taking the $\arg \max$ of attention scores α_i in a zero-shot manner.

2.4 Results

We show the results in Table 1. The performance of our approach is comparable to previous work. Note that our model does not use any heuristics to prune out utterances. Instead, the attention scores α_i are decided entirely by the model. We also run an experiment using augmented data following Wu et al. (2020), which is constructed by excerpting partial context from the original context². Finally, we want to highlight the importance of the attention module \mathcal{A} , where the performance drops by 10 points if removed.

2.5 Attention Analysis

The empirical success of the hierarchical encoder has an important implication: it is able to link the candidate with one or multiple relevant utterances in the context. This can be proved by the attention distribution α_i . Intuitively, if C_k is the correct next response (i.e. $k = 1$), then the attention distribution should be sharp, which indicates an implicit *reply-to* that links to one of the previous utterances. In contrast, if C_k is incorrect (i.e. $k \neq 1$), our model is less likely to find an implicit link, and the attention distribution should be flat. Entropy is a good tool

²For a context of length 50, we can take the first $i \in 1 \dots 49$ utterances as a new context and the $i + 1$ -th utterance acts as the new correct next response. We can sample the negatives randomly from other sessions in the dataset.

to quantify sharpness. Numerically, the entropy is 1.4 (sharp) when C_k is correct and 2.1 (flat) for incorrect ones, validating our suppositions.

Is it possible to reveal these implicit links? The solution is inspired by the labeled data of dialogue disentanglement as elaborated in §3.4.

3 Zero shot Dialogue disentanglement

3.1 Task Description

The dataset used is DSTC8 subtask-4 (Kummerfeld et al., 2018)³. We want to find the parent utterance in an entangled context to which the current utterance is replying, and repeat this process for every utterance. After all the links are predicted, we run a connected component algorithm over them, where each connected component is one thread.

3.2 Related Work

All previous work (Shen et al., 2006; Elsner and Charniak, 2008; Wang and Oard, 2009; Elsner and Charniak, 2011; Jiang et al., 2018; Kummerfeld et al., 2018; Zhu et al., 2020; Li et al., 2020; Yu and Joty, 2020) treat the task as a sequence of multiple-choice problems. Each of them consists of a sliding window of n utterances. The task is to link the last

³It is a disjoint split of the Ubuntu IRC channel as opposed to the one used in §2.1, hence there is no risk of information leak.

Setting	w	data%	CLUSTER					LINK		
			VI	ARI	P	R	F1	P	R	F1
1) zero shot										
	0.00	0.0	62.9	14.7	2.9	0.3	0.5	41.2	39.7	40.5
	0.25	0.0	84.4	50.1	25.9	24.8	25.3	43.7	41.4	42.2
	0.50	0.0	84.6	51.5	24.6	23.8	24.2	41.5	40.0	40.8
	0.75	0.0	84.6	49.2	23.3	23.1	23.2	41.8	40.3	41.1
	1.00	0.0	84.3	47.5	22.9	23.0	23.0	41.6	40.1	40.9
2) few shot										
finetune	0.25	1	89.7	60.2	26.1	33.9	29.5	65.0	62.7	63.8
scratch	-	-	88.7	58.7	22.6	28.6	25.2	63.7	61.4	62.6
finetune	0.25	10	90.6	59.8	32.4	38.4	35.1	70.5	68.0	69.3
scratch	-	-	90.4	61.0	32.4	36.5	34.3	70.4	67.9	69.1
finetune	0.25	100	91.1	62.7	35.3	42.0	38.3	74.2	71.6	72.9
scratch	-	-	91.2	62.1	35.6	40.3	37.8	74.0	71.3	72.6

Table 2: $w = 0$ indicates pure entangled response selection training. In the few-shot section, scratch is the disentanglement model not trained self-supervisedly on entangled response selection before. The evaluation metrics and labeled data used for fine-tuning are in Kummerfeld et al. (2018). Results are the average of three runs.

utterance to one of the preceding $n - 1$ utterances. This model is usually trained in supervised mode using the labeled *reply-to* links. Our model also follows the same formulation.

3.3 Model (Dashed Arrows in Figure 2)

We use the trained hierarchical model in §2.3 without the final MLP layer used for scoring. In addition, we only have one candidate now, which is the last utterance in a dialogue. We use $C_{k=1}$ to represent it for consistency. Note that we only need to calculate $i' = \arg \max_i \alpha_i$. This indicates that $C_{k=1}$ is replying to utterance $U_{i'}$ in the context.

3.4 Proposed Attention Supervision

We note that the labeled *reply-to* links act as supervision to the attention α_i : they indicate which α_i should be 1. We call this extrinsic supervision. Recall the implicit attention analysis in §2.5, from which we exploit two kinds of intrinsic supervision:

- If C_k is the correct next response, then $\alpha_{n+1} = 0$ because C_k should be linking to one previous utterance, not itself.
- If C_k is incorrect, then it should point to itself, acting like the start utterance of a new thread. Hence, $\alpha_{n+1} = 1$.

We train this intrinsic attention using MSE (row 2 of the multi-task loss table in Figure 2) along with the original response selection loss using a weight w for linear combination $L = (1 - w) * L_{res} + w * L_{attn}$. Note that we do not use any labeled

disentanglement data in the training process.

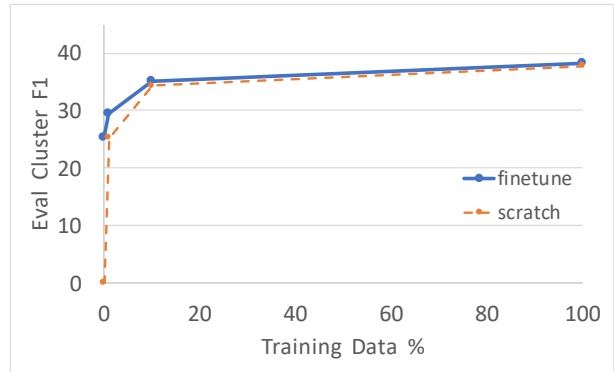


Figure 3: Different amounts of labeled data for fine-tuning. The model with self-supervised response selection training outperforms the one trained from scratch.

3.5 Results

We present the results in Table 2. In the first section, we focus on zero-shot performance, where we vary w to see its effect. As we can see, $w = 0.25$ gives a close-to-best performance in terms of cluster and link scores. Therefore, we use it for few-shot fine-tuning setting, under which our proposed method outperforms baselines trained from scratch by a large margin. We pick the best checkpoint based on the validation set performance and evaluate it on the test set. This procedure is repeated three times with different random seeds to get the averaged performance reported in Table 2. With 10% of the data, we can achieve 92% of the performance

trained using full data. The performance gap becomes smaller when more data is used as illustrated in Figure 3.

3.6 Real-World Application

Our method only requires one additional MLP layer attached to the architecture of Li et al. (2020) to train on the entangled response selection task, hence it is trivial to swap the trained model into a production environment. Suppose a dialogue disentanglement system (Li et al., 2020) is already up and running:

1. Train a BERT model on the entangled response selection task (§2.1) with attention supervision loss (§3.4). This is also the multi-task loss depicted in Figure 2.
2. Copy the weight of the pretrained model into the existing architecture (Li et al., 2020).
3. Perform zero-shot dialogue disentanglement (zero-shot section of Table 2) right away, or finetune the model further when more labeled data becomes available (few-shot section of Table 2).

This strategy will be useful especially when we want to bootstrap a system with limited and expensive labeled data.

4 Conclusion

In this paper, we first demonstrate that entangled response selection does not require complex heuristics for context pruning. This implies the model might have learned implicit *reply-to* links useful for dialogue disentanglement. By introducing intrinsic attention supervision to shape the distribution, our proposed method can perform zero-shot dialogue disentanglement. Finally, with only 10% of the data for tuning, our model can achieve 92% of the performance of the model trained on full labeled data. Our method is the first attempt to zero-shot dialogue disentanglement, and it can be of high practical value for real-world applications.

References

Iz Beltagy, Matthew E Peters, and Arman Cohan. 2020. Longformer: The long-document transformer. *arXiv preprint arXiv:2004.05150*.

Dario Bertero, Kenichi Yokote Takeshi Homma and, Makoto Iwayama, and Kenji Nagamatsu. 2020.

Model ensembling of esim and bert for dialogue response selection. *DSTC8*.

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*.

Micha Elsner and Eugene Charniak. 2008. You talking to me? a corpus and algorithm for conversation disentanglement. In *Proceedings of ACL-08: HLT*, pages 834–842.

Micha Elsner and Eugene Charniak. 2011. Disentangling chat with local coherence models. In *Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics: Human Language Technologies*, pages 1179–1189.

Jia-Chen Gu, Tianda Li, Quan Liu, Xiaodan Zhu, Zhen-Hua Ling, and Yu-Ping Ruan. 2020. Pre-trained and attention-based neural networks for building noetic task-oriented dialogue systems. *arXiv preprint arXiv:2004.01940*.

Chulaka Gunasekara, Jonathan K. Kummerfeld, Luis Lastras, and Walter S. Lasecki. 2020. Noesis ii: Predicting responses, identifying success, and managing complexity in task-oriented dialogue. In *8th Edition of the Dialog System Technology Challenges at AAAI 2019*.

Jyun-Yu Jiang, Francine Chen, Yan-Ying Chen, and Wei Wang. 2018. Learning to disentangle interleaved conversational threads with a siamese hierarchical network and similarity ranking. 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 1812–1822.

Jonathan K Kummerfeld, Sai R Gouravajhala, Joseph Peper, Vignesh Athreya, Chulaka Gunasekara, Jatin Ganhotra, Siva Sankalp Patel, Lazaros Polymenakos, and Walter S Lasecki. 2018. A large-scale corpus for conversation disentanglement. *arXiv preprint arXiv:1810.11118*.

Tianda Li, Jia-Chen Gu, Xiaodan Zhu, Quan Liu, Zhen-Hua Ling, Zhiming Su, and Si Wei. 2020. Dialbert: A hierarchical pre-trained model for conversation disentanglement. *arXiv preprint arXiv:2004.03760*.

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

Dou Shen, Qiang Yang, Jian-Tao Sun, and Zheng Chen. 2006. Thread detection in dynamic text message streams. In *Proceedings of the 29th annual international ACM SIGIR conference on Research and development in information retrieval*, pages 35–42.

- Lidan Wang and Douglas W Oard. 2009. Context-based message expansion for disentanglement of interleaved text conversations. In *Proceedings of human language technologies: The 2009 annual conference of the North American chapter of the association for computational linguistics*, pages 200–208. Citeseer.
- Weishi Wang, Shafiq Joty, and Steven CH Hoi. 2020. Response selection for multi-party conversations with dynamic topic tracking. *arXiv preprint arXiv:2010.07785*.
- Shuangzhi Wu, Xu Wang Yufan Jiang and, Wei Mia, Zhenyu Zhao, Jun Xie, and Mu Li. 2020. Enhancing response selection with advanced context modeling and post-training. *DSTC8*.
- Zhilin Yang, Zihang Dai, Yiming Yang, Jaime Carbonell, Ruslan Salakhutdinov, and Quoc V Le. 2019. Xlnet: Generalized autoregressive pretraining for language understanding. *arXiv preprint arXiv:1906.08237*.
- Tao Yu and Shafiq Joty. 2020. Online conversation disentanglement with pointer networks. *arXiv preprint arXiv:2010.11080*.
- Manzil Zaheer, Guru Guruganesh, Avinava Dubey, Joshua Ainslie, Chris Alberti, Santiago Ontanon, Philip Pham, Anirudh Ravula, Qifan Wang, Li Yang, et al. 2020. Big bird: Transformers for longer sequences. *arXiv preprint arXiv:2007.14062*.
- Henghui Zhu, Feng Nan, Zhiguo Wang, Ramesh Nallapati, and Bing Xiang. 2020. Who did they respond to? conversation structure modeling using masked hierarchical transformer. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, pages 9741–9748.