

Unsupervised Conversation Disentanglement through Co-Training

Hui Liu Zhan Shi Xiaodan Zhu

Ingenuity Labs Research Institute & ECE, Queen’s University, Canada

{hui.liu, z.shi, xiaodan.zhu}@queensu.ca

Abstract

Conversation disentanglement aims to separate intermingled messages into detached sessions, which is a fundamental task in understanding multi-party conversations. Existing work on conversation disentanglement relies heavily upon human-annotated datasets, which are expensive to obtain in practice. In this work, we explore to train a conversation disentanglement model without referencing any human annotations. Our method is built upon a deep co-training algorithm, which consists of two neural networks: a message-pair classifier and a session classifier. The former is responsible for retrieving local relations between two messages while the latter categorizes a message to a session by capturing context-aware information. Both networks are initialized respectively with pseudo data built from an unannotated corpus. During the deep co-training process, we use the session classifier as a reinforcement learning component to learn a session assigning policy by maximizing the local rewards given by the message-pair classifier. For the message-pair classifier, we enrich its training data by retrieving message pairs with high confidence from the disentangled sessions predicted by the session classifier. Experimental results on the large Movie Dialogue Dataset demonstrate that our proposed approach achieves competitive performance compared to the previous supervised methods. Further experiments show that the predicted disentangled conversations can promote the performance on the downstream task of multi-party response selection.

1 Introduction

With the continuing growth of Internet and social media, online group chat channels, e.g., Slack¹ and Whatsapp², among many others, have become increasingly popular and played a significant social

¹<https://slack.com/>

²<https://www.whatsapp.com/>







Session	Message
S1	 Anyone finished the assignment?
S1	 Not yet. Working on it
S2	 I'd like to get a new keyboard. Any suggestions?
S1	 When is the due date?
S2	 The one I have is pretty good.
S1	 Hmm... tonight?

Figure 1: An example of conversation disentanglement with two sessions occurring at the same time.

and economic role. Along with the convenience of instant communication brought by these applications, the inherent property that multiple topics are often discussed in one channel hinders an efficient access to the conversational content. In the example shown in Figure 1, people or intelligent systems have to selectively read the messages related to the topics they are interested in from hundreds of messages in the chat channel.

With the goal of automatically grouping messages with the same topic into one session, *conversation disentanglement* has proved to be a prerequisite for understanding multi-party conversations and solving the corresponding downstream tasks such as response selection (Elsner and Charniak, 2008; Lowe et al., 2017; Jia et al., 2020; Wang et al., 2020). Previous research on conversation disentanglement can be roughly divided into two categories: (1) two-step methods, and (2) end-to-end methods. In the two-step methods (Elsner and Charniak, 2011, 2008; Jiang et al., 2018), a model first retrieves the “local” relations between two messages by utilizing either feature engineering approaches or deep learning methods, and then a clustering algorithm is employed to divide an entire conversation into separate sessions based on the message pair relations. In contrast, end-to-end methods (Tan et al., 2019; Yu and Joty, 2020) capture the “global” information contained in the context of detached




Figure 2: Illustration of our proposed co-training framework.

sessions and calculate the matching degree between a session and a message in an end-to-end manner. Though end-to-end methods have been proved to be more flexible and can achieve better performance (Liu et al., 2020), these two types of methods are interconnected and complementary since a global optimal clustering solution on the local relations will produce the optimal disentanglement scheme (McCallum and Wellner, 2004).

Although the previous research efforts have achieved an impressive progress on conversation disentanglement, they all highly rely on human-annotated corpora, which are expensive and scarce to obtain in practice (Kummerfeld et al., 2019). The heavy dependence on human annotations limits the extensions of related study on conversation disentanglement as well as the applications on downstream tasks, given a wide variety of occasions where multi-party conversations can happen. In this work, we explore the possibility to train an end-to-end conversation disentanglement model without referencing any human annotations and propose a completely unsupervised disentanglement model.

Our method builds upon the co-training approach (Blum and Mitchell, 1998; Nigam and Ghani, 2000) but extends it to a deep learning framework. By viewing the disentanglement task from the local perspective and the global perspective, our method consists of a message-pair classifier and a session classifier. The message-pair classifier aims to retrieve the message pair relations, which is of a similar purpose as the model used in a two-step method that retrieves the local relations between two messages. The session classifier is a global context-aware model that can directly categorize a message into a session in an end-to-end fashion. The two classifiers view the task of conversation disentanglement from the perspectives of a local two-step method and an global end-to-end

model, which will be separately initialized with pseudo data built from the unannotated corpus and updated with each other during co-training. More concretely, during the co-training procedure, we adopt reinforcement learning to learn a session assigning policy for the session classifier by maximizing the accumulated rewards between a message and a session which are given by the message-pair classifier. After updating the parameters of the session classifier, a new set of data with high confidence will be retrieved from the predicted disentanglement results of the session classifier and used for updating the message-pair classifier. As shown in Figure 2, the above process is iteratively performed by updating one classifier with the other until the performance of session classifier stops increasing.

We conduct experiments on the large public Movie Dialogue Dataset (Liu et al., 2020). Experimental results demonstrate that our proposed method outperforms strong baselines based on BERT (Devlin et al., 2019) in two-step settings, and achieves competitive results compared to those of the state-of-the-art supervised end-to-end methods. Moreover, we apply the disentangled conversations predicted by our method to the downstream task of multi-party response selection and get significant improvements compared to a baseline system.³ In summary, our main contributions are three-fold:

- To the best of our knowledge, this is the first work to investigate unsupervised conversation disentanglement with deep neural models.
- We propose a novel approach based on co-training which can perform unsupervised conversation disentanglement in an end-to-end fashion.
- We show that our method can achieve performance competitive with supervised methods on the large public Movie Dialogue Dataset. Further experiments show that our method can be easily adapted to downstream tasks and achieve significant improvements.

2 Related Work

Conversation Disentanglement Conversation disentanglement has long been regarded as a fundamental task for understanding multi-party conversations (Elsner and Charniak, 2008, 2010) and

³Code will be publicly available at https://github.com/LayneIns/Unsupervised_dialo_disentanglement

can be combined with downstream tasks to boost their performance (Jia et al., 2020; Wang et al., 2020). Previous methods on conversation disentanglement are mostly performed in a supervised fashion, which can be classified as two categories: (1) two-step approaches and (2) end-to-end methods. The two-step methods (Elsner and Charniak, 2008, 2010, 2011; Chen et al., 2017; Jiang et al., 2018; Kummerfeld et al., 2019) firstly retrieve the relations between two messages, e.g., “reply-to” relations (Guo et al., 2018; Zhu et al., 2020), and then adopt a clustering algorithm to construct individual sessions. The end-to-end models (Tan et al., 2019; Liu et al., 2020; Yu and Joty, 2020), instead, perform the disentanglement operation in an end-to-end manner, where the context information of detached sessions will be exploited to classify a message to a session. End-to-end models tend to achieve better performance than two-step models, but both often need large annotated data to get fully trained (Liu et al., 2020), which is expensive to obtain and thus encourages the demand on unsupervised algorithms. A few preliminary studies perform unsupervised thread detection in email systems based on two-step methods (Wu and Oard, 2005; Erera and Carmel, 2008; Domeniconi et al., 2016), but these methods use handcrafted features which cannot be extended to various datasets. Compared with previous works, our method can conduct end-to-end conversation disentanglement in a completely unsupervised fashion, which can be easily adapted to downstream tasks and used in a wide variety of applications.

Dialogue Structure Learning One problem that may be related to conversation disentanglement is dialogue structure learning (Zhai and Williams, 2014; Shi et al., 2019). Both are related to understanding multi-party conversation structures but they are different tasks. Dialogue structure learning aims to discover latent dialogue topics and construct an implicit utterance dependency tree to represent a multi-party dialogue’s turn taking (Qiu et al., 2020), while the goal of conversation disentanglement is to learn an explicit dividing scheme that separates intermingled messages into sessions.

Co-training Co-training (Blum and Mitchell, 1998; Nigam and Ghani, 2000) has been widely used as a low-resource learning algorithm in natural language processing (Wu et al., 2018; Chen et al., 2018), which assumes that the data has two

complementary views and utilizes two models to iteratively provide pseudo training signals to each other. Our method consists of a message-pair classifier and a session classifier, which respectively view the unannotated dataset from the perspective of the local relations between two messages and that of the context-aware relations between a session and a message. To the best of our knowledge, this is the first work that utilizes co-training in the research of conversation disentanglement. We will extend the co-training idea to the deep learning paradigm to construct novel models for disentanglement.

3 Formulation and Notations

Given a conversation $\mathcal{C} = [m_1, m_2, \dots, m_N]$ where m_i is a message with speaker \mathcal{S}_i , it contains K sessions $\{T^k\}_{k=1}^K$ where $T^k = [m_1^k, \dots, m_{|T^k|}^k]$ and K are unknown to the model. Our goal is to learn a dividing scheme that indicates which session a message m_i belongs to. We solve this task in an end-to-end fashion where we formulate unsupervised conversation disentanglement as an unsupervised sequence labeling task. For a given message m_i , there exists a session set $\mathbb{T} = \{T^1, \dots, T^{z(i)}\}$ where $z(i)$ indicates the number of detached sessions when m_i is being processed. The model needs to decide if m_i belongs to any session in \mathbb{T} . If $m_i \in T^k$, then m_i will be appended to T^k ; otherwise a new session $T^{z(i)+1}$ will be built and initialized by m_i , and the new session will be added to \mathbb{T} .

4 Method

In this section, we describe our co-training based framework in detail, which contains following components:

1. A message-pair classifier which can retrieve the relations between two messages. The relation scores will be used as rewards for updating the session classifier during co-training.
2. A session classifier which can perform end-to-end conversation disentanglement by retrieving the relations between a message and a session. The predicted results will be used to build new pseudo data to train the message-pair classifier during co-training.
3. A co-training algorithm involving the message-pair classifier and the session classifier. Two classifiers will help to update

each other until the performance of the session classifier stops growing.

We will introduce the details of the three components in following sections.

4.1 Message-pair Classifier

The message-pair classifier is a binary classifier which we denote as \mathcal{F}_m in the remainder of this paper. Due to the lack of annotated data in unsupervised settings, the goal of \mathcal{F}_m is to predict if two messages are in the same session; i.e., whether they talk about the same topic, which is different from most previous work that predicts the “reply-to” relation. In our experiment, we adopt a pretrained BERT (Devlin et al., 2019) in the base version as our message encoder.

4.1.1 Model

Given two messages m_i and m_j , we separately obtain the sentence embeddings of the two messages:

$$v_{m_i} = \text{BERT}(m_i) \quad (1)$$

$$v_{m_j} = \text{BERT}(m_j) \quad (2)$$

The probability of m_i and m_j belonging to the same session is computed with the dot product between v_{m_i} and v_{m_j} :

$$p_m = \text{Sigmoid}(v_{m_i} \cdot v_{m_j}) \quad (3)$$

We abbreviate Eq. 1–3 as:

$$p_m(m_i, m_j) = \mathcal{F}_m(m_i, m_j) \quad (4)$$

\mathcal{F}_m is trained to minimize the cross-entropy loss. The predicted probabilities between message pairs will be used as rewards during the co-training process to update the session classifier.

4.1.2 Initialization

One important step in standard co-training algorithm is to initialize the classifiers with a small amount of annotated data. Since our dataset is completely unlabeled, we create a pseudo dataset to initialize the message-pair classifier. The assumption we use in our experiments is that one speaker mostly participates in only one session⁴.

⁴We verify this assumption in two natural multi-party conversation datasets: Reddit dataset (Tan et al., 2019) and Ubuntu IRC dataset (Kummerfeld et al., 2019). Statistics show that only 6% of speakers will join multiple sessions on the Reddit dataset and 20% on the IRC dataset.

To construct the pseudo data \mathcal{D}_m , we use the message pairs from the same speaker in one conversation as the positive cases, while randomly sampling messages from different conversations as the negative pairs. In this way we obtain a retrieved dataset \mathcal{D}_m^{ret} containing 937K positive cases and 2,184K negative cases. However, we observe that the positive cases constructed from the above process are very noisy because: (1) there are still some speakers who will appear in multiple sessions, and (2) even message pairs from the same speaker in the same session can be very semantically different since they are not contiguous messages. These noisy training cases will result in low confidences for the predicted probabilities of \mathcal{F}_m , which will be used later in co-training. Thus we randomly select some messages from the unlabeled dataset and use a pretrained DialoGPT (Zhang et al., 2020) to generate direct responses in order to form new positive cases, which we denote as \mathcal{D}_m^{gen} . In this way, we finally obtain the pseudo data $\mathcal{D}_m = \mathcal{D}_m^{ret} \cup \mathcal{D}_m^{gen}$, which contains 1,212K positive cases and 2,184K negative cases, to initialize \mathcal{F}_m .

4.1.3 Two-step Disentanglement

After being trained on the pseudo data \mathcal{D}_m , the message-pair classifier \mathcal{F}_m can be exploited as for two-step conversation disentanglement. Given an unlabeled conversation as $\mathcal{C} = [m_1, m_2, \dots, m_N]$, we first use \mathcal{F}_m to predict the probability between each message pair in \mathcal{C} . Then we perform the greedy search algorithm widely used in the previous works (Elsner and Charniak, 2008) to segment \mathcal{C} into detached sessions.

4.2 Session Classifier

The session classifier, denoted as \mathcal{F}_t , aims to calculate the relations between a session and a message that indicates if the message belongs to the session or not. Given the current context of a session as $T = [m_1, \dots, m_{|T|}]$ and a message m , the goal of \mathcal{F}_t is to decide if m can be appended to T or not.

4.2.1 Model

For each message $m_j \in T$, we obtain its sentence embedding v_{m_j} by a Bidirectional LSTM network (Hochreiter and Schmidhuber, 1997) and a multi-layer perceptron (MLP):

$$\vec{v}_{m_j}, \overleftarrow{v}_{m_j} = \text{BiLSTM}(m_j) \quad (5)$$

$$v_{m_j} = \text{MLP}([\vec{v}_{m_j}, \overleftarrow{v}_{m_j}]) \quad (6)$$

After obtaining sentence embeddings of all the messages in T as $[v_{m_1}, \dots, v_{m_{|T|}}]$, we adopt a self attention mechanism (Yang et al., 2016) to calculate the session embedding v_T by aggregating the information from different messages. Specifically,

$$u_{m_j} = \tanh(w \cdot v_{m_j} + b) \quad (7)$$

$$\alpha_{m_j} = \frac{\exp(u_{m_j})}{\sum_j \exp(u_{m_j})} \quad (8)$$

$$v_T = \sum_j \alpha_{m_j} v_{m_j} \quad (9)$$

where w and b are trainable parameters.

For the message m , we use the same Bidirectional LSTM network and MLP as in Equation 5 and 6 to obtain its sentence embedding v_m . Then the probability of m belonging to T is calculated with the dot product between v_m and v_T :

$$p_t = \text{Sigmoid}(v_T \cdot v_m) \quad (10)$$

We abbreviate the above process as:

$$p_t(T, m) = \mathcal{F}_t(v_T, v_m) \quad (11)$$

\mathcal{F}_t is trained to minimize the cross-entropy loss.

4.2.2 Initialization

Similar to the message-pair classifier, we build a pseudo dataset \mathcal{D}_t to initialize the session classifier \mathcal{F}_t to make it be able to decide if a message is semantically consistent with a sequence of messages. We construct \mathcal{D}_t based on the same assumption that one speaker is involved in just one session for most of the time.

Given a conversation $\mathcal{C} = [m_1, m_2, \dots, m_N]$ from the unlabeled corpus, we retrieve the messages from a speaker \mathcal{S} as $\mathcal{C}_{\mathcal{S}} = [m_{\mathcal{S}_1}, m_{\mathcal{S}_2}, \dots]$ where $\mathcal{C}_{\mathcal{S}} \subset \mathcal{C}$. Based on the assumption, the messages in $\mathcal{C}_{\mathcal{S}}$ are in the same session, so the message $m_{\mathcal{S}_i} \in \mathcal{C}_{\mathcal{S}}$ where $i \neq 1$ and its preceding context can be regarded as the positive input of \mathcal{F}_t . Consider a positive case with $m_{\mathcal{S}_2}$ as the message, the message m and session T are defined as follows:

$$m := m_{\mathcal{S}_2} \quad (12)$$

$$T := [m_1, \dots, m_{\mathcal{S}_2-1}] \quad (13)$$

The reason is that $m_{\mathcal{S}_1} \in [m_1, \dots, m_{\mathcal{S}_2-1}]$, so $[m_1, \dots, m_{\mathcal{S}_2-1}]$ and $m_{\mathcal{S}_2}$ should be semantically consistent according to the assumption.

For the negative instances of \mathcal{D}_t , we randomly sample a conversation as T from the corpus, and a message from another conversation as m . As such we obtain a pseudo dataset \mathcal{D}_t consisting of 460K positive instances and 1,158K negative cases.

Algorithm 1 An end-to-end method for conversation disentanglement with the session classifier

Input: An unlabeled conversation \mathcal{C} , the initialized session classifier \mathcal{F}_t

Output: A set of sessions \mathbb{T}

```

1: Let  $\mathbb{T} = \emptyset$ 
2: for  $m_i \in \mathcal{C}$  do
3:   if  $\mathbb{T}$  is empty then
4:     Let  $T^{z(i)+1} = \{m_i\}$ 
5:      $\mathbb{T} = \mathbb{T} \cup \{T^{z(i)+1}\}$ 
6:   else
7:     Let  $\mathcal{C}_i = [m_1, \dots, m_{i-1}]$ 
8:      $p_t(\mathcal{C}_i, m_i) = \mathcal{F}_t(v_{\mathcal{C}_i}, v_{m_i})$ 
9:     if  $p_t(\mathcal{C}_i, m_i) < 0.5$  then
10:      Let  $T^{z(i)+1} = \{m_i\}$ 
11:       $\mathbb{T} = \mathbb{T} \cup \{T^{z(i)+1}\}$ 
12:    else
13:      Let  $P_m(i) = \{\}$ 
14:      for  $T^k \in \mathbb{T}$  do
15:         $p_t(T^k, m_i) = \mathcal{F}_t(v_{T^k}, v_{m_i})$ 
16:        Add  $p_t(T^k, m_i)$  to  $P_m(i)$ 
17:      end for
18:       $T_{\max}^k = \arg \max P_m(i)$ 
19:       $T_{\max}^k = T_{\max}^k \cup \{m_i\}$ 
20:    end if
21:  end if
22: end for

```

Return \mathbb{T}

4.2.3 End-to-end Disentanglement

Note that after initialized with the pseudo data \mathcal{D}_t , the session classifier \mathcal{F}_t can be directly applied to perform end-to-end conversation disentanglement. Suppose message m_i is being processed where $m_i \in \mathcal{C}$ and $\mathcal{C} = [m_1, m_2, \dots, m_N]$, we first calculate the probability of m_i belonging to its preceding context $\mathcal{C}_i = [m_1, \dots, m_{i-1}]$, which we denote as $p_t(\mathcal{C}_i, m_i) = \mathcal{F}_t(v_{\mathcal{C}_i}, v_{m_i})$. If $p_t(\mathcal{C}_i, m_i) < 0.5$, m_i will be used to initialize a new session $T^{z(i)+1}$ where z is a function indicating the number of disentangled sessions in \mathcal{C}_i ; otherwise m_i will be used to calculate the matching probability with each session in \mathbb{T} , and then be classified to the session which has the greatest matching probability. The overall end-to-end algorithm is shown in Algorithm 1.

4.3 Co-Training

The confidence of \mathcal{F}_m and \mathcal{F}_t is not high because they are initialized with noisy pseudo data. We

propose to adapt the idea of co-training to the disentanglement task, which is leveraged to iteratively update the two classifiers with the help of each other. The session classifier will utilize the local probability provided by the message-pair classifier with reinforcement learning, while more training data, built from the outcomes of the session classifier, will be fed to the message-pair classifier. We will introduce the details in this subsection.

4.3.1 Updating Session Classifier

Since no labeled data is provided to train \mathcal{F}_t , we formulate the disentanglement task as a deterministic Markov Decision Process and adopt the Policy Gradient algorithm (Sutton et al., 1999) for the optimization. For each co-training iteration, \mathcal{F}_t will be initialized with the pseudo data \mathcal{D}_t and then updated by reinforcement learning.

State The state s_i of the i^{th} disentanglement step consists of three components (m_i, C_i, \mathbb{T}) , where m_i is the i^{th} message of C ; $C_i = [m_1, \dots, m_{i-1}]$ is the preceding context of m_i ; \mathbb{T} is the detached session set which contains $z(i)$ sessions.

Action The action space of the i^{th} disentanglement step consists of two types of actions:

1. Classifying m_i to a new session, which we denote as $a_i^{\text{new}} \in \{0, 1\}$. If a_i^{new} is 0, m_i will be used to initialize a new session $T^{z(i)+1}$; otherwise m_i will be categorized into an existing session.
2. Categorizing m_i to an existing session in \mathbb{T} , which we denote as $a_i^t \in \{1, \dots, z(i)\}$.

Policy network We parameterize the action with a policy network π which is in a hierarchical structure. The first layer policy $\pi^{\text{new}}(\mathbf{a}_i^{\text{new}} | s_i; \theta_1)$ is to decide if a message m_i belongs to C_i , and the first layer action $a_i^{\text{new}} \in \{0, 1\}$ will be sampled:

$$\pi^{\text{new}}(\mathbf{a}_i^{\text{new}} | s_i; \theta_1) = p_t(C_i, m_i) \quad (14)$$

$$a_i^{\text{new}} \sim \pi^{\text{new}}(\mathbf{a}_i^{\text{new}} | s_i; \theta_1) \quad (15)$$

If a_i^{new} is 1, which means m_i belongs to a session in \mathbb{T} , the second layer policy $\pi^t(\mathbf{a}_i^t | s_i; \theta_2)$ will decide which of existing sessions that m_i should be categorized to:

$$\pi^t(\mathbf{a}_i^t | s_i; \theta_2) = \{p_t(T^k, m_i) | T^k \in \mathbb{T}\} \quad (16)$$

$$a_i^t \sim \pi^t(\mathbf{a}_i^t | s_i; \theta_2) \quad (17)$$

where θ_1 and θ_2 are both parameters.

Reward The rewards are provided by the message-pair classifier \mathcal{F}_m . For $a_i^{\text{new}} = 0$, we want m_i to be different from all the messages in C_i . Thus it is defined by the negative average of the probabilities between m_i and all the messages in C_i . However, for $a_i^{\text{new}} = 1$ and $a_i^t = k$, we want m_i to be similar to all the messages in T^k , and thus the reward is defined as the average of the probabilities between m_i and all the messages in T^k :

$$r_{m_i} = \begin{cases} -\text{Avg}(\{\mathcal{F}_m(m_i, m_j) | m_j \in C_i\}), & a_i^{\text{new}} = 0 \\ \text{Avg}(\{\mathcal{F}_m(m_i, m_j) | m_j \in T^k\}), & a_i^t = k \end{cases} \quad (18)$$

An issue associated with r_{m_i} is that its confidence might be low because \mathcal{F}_m is trained on noisy pseudo data. We hence design another speaker reward r_{s_i} based on our assumptions. For a message m_i initializing a new session $T^{z(i)+1}$, its speaker s_i should not appear in C_i ; while for a message m_i categorized to an existing session T^k , it should receive a positive reward if its speaker s_i appears in T^k :

$$r_{s_i} = \begin{cases} -1, & a_i^{\text{new}} = 0 \text{ and } s_i \in C_i \\ 1, & a_i^t = k \text{ and } s_i \in T^k \\ 0, & \text{otherwise} \end{cases} \quad (19)$$

The final reward r_i for an action is calculated as:

$$r_i = \gamma r_{m_i} + (1 - \gamma) r_{s_i} \quad (20)$$

where γ is a parameter ranged in $[0, 1]$ that balances r_{m_i} and r_{s_i} , which we set to 0.6 in experiments. The policy network parameters θ_1 and θ_2 are learned by optimizing:

$$J(\theta_1, \theta_2) = E_{(\pi^{\text{new}}, \pi^t)} \left[\sum_{i=1}^N r_i \right] \quad (21)$$

4.3.2 Updating Message-pair Classifier

As mentioned in Section 4.1.2, the pseudo data \mathcal{D}_m for initializing the message-pair classifier \mathcal{F}_m is noisy. Thus we enrich \mathcal{D}_m with new training instances $\mathcal{D}_m^{\text{new}}$ retrieved from the predicted disentanglement results of \mathcal{F}_t .

Given a conversation C , \mathcal{F}_t can predict the disentangled sessions as $\mathbb{T} = \{T^k\}_{k=1}^K$. Given session $T^k = [m_1^k, \dots, m_{|T^k|}^k]$ as an example, for a message m_i^k , we retrieve its

Method Type	Method	NMI	1-1	Loc ₃	Shen-F	MSE↓
Two-step	Vanilla BERT	0.0	40.20	49.26	50.50	1.9710
	BERT + \mathcal{D}_m^{ret}	10.59	45.62	53.4	53.47	1.4750
	BERT + \mathcal{D}_m	11.13	45.74	53.69	53.64	1.4617
	BERT + \mathcal{D}_m + CT	11.32	45.89	53.80	53.75	1.4602
	*BERT	11.52	45.99	54.04	53.87	1.4636
End-to-end	Session Classifier + \mathcal{D}_t	24.96	54.26	60.66	59.16	0.8059
	Session Classifier + \mathcal{D}_t + CT	29.71	56.38	62.46	60.44	0.6871
	*Liu et al. (2020)	35.30	57.31	63.27	64.37	0.5299

Table 1: The results of conversation disentanglement. * means the method is performed in a supervised manner. “CT” represents co-training. “Vanilla BERT” represents BERT without finetuning. Note that $\mathcal{D}_m^{ret} \cup \mathcal{D}_m^{gen} = \mathcal{D}_m$, where \mathcal{D}_m^{ret} is the pairs retrieved from the unannotated corpus and \mathcal{D}_m^{gen} is generated by DialoGPT. ↓ means the lower the better.

preceding M messages in T^k and form M pairs $\{(m_{i-M}^k, m_i^k), (m_{i-(M-1)}^k, m_i^k), \dots, (m_{i-1}^k, m_i^k)\}$ as the new positive pseudo message pairs. In order to raise the confidence of the newly added data, we filter out those pairs where the two messages have less than 2 overlapped tokens after removing stopwords. For each co-training iteration, \mathcal{F}_m is retrained on the data $\mathcal{D}_m \cup \mathcal{D}_m^{new}$.

5 Experiments

5.1 Experimental Setup

5.1.1 Dataset

A large corpus is often required for end-to-end conversation disentanglement. In this work, we conduct experiments on the publicly available Movie Dialogue Dataset (Liu et al., 2020) which is built from online movie scripts. It contains 29,669/2,036/2,010 instances for train/dev/test split with a total of 827,193 messages, where the session number in one instance can be 2, 3 or 4. Since we make explorations in unsupervised settings, no labels are used in our training.

5.1.2 Implementation Details

We adopt BERT (Devlin et al., 2019) (the uncased base version) as the message-pair classifier. For the session classifier, we set the hidden dimension to be 300, and the word embeddings are initialized with 300-d GloVe vectors (Pennington et al., 2014). For training, we use Adam (Kingma and Ba, 2015) for optimization; the learning rate is set to be 1e-5 for the message-pair classifier, 1e-4 for initializing the session classifier, and 1e-5 for updating the session classifier with reinforcement learning. We

iterate for 3 turns for co-training when the best performance is achieved on the development set.

5.1.3 Evaluation Metrics

Four clustering metrics widely used in the previous work (Elsner and Charniak, 2008; Kummerfeld et al., 2019; Tan et al., 2019) are adopted: Normalized mutual information (NMI), One-to-One Overlap (1-1), Loc₃ and Shen F score (Shen-F). More explanations about the metrics can be found in Appendix A.1.

We also report the mean squared error (MSE) between the predicted session numbers and the golden session numbers as previous work (Liu et al., 2020). This metric can measure whether the model can disentangle a given dialogue to the correct number of sessions.

5.2 Results

5.2.1 Disentanglement Performance

Table 1 shows the results of unsupervised conversation disentanglement of different methods. We can observe that: (1) for two-step methods, BERT has a very poor performance without finetuning, while after finetuned on our pseudo dataset, its performance gets improved with a relatively large margin. (2) Utilizing the pseudo pairs generated by a pretrained DialoGPT can further improve the performance of BERT based on \mathcal{D}_m^{ret} . We consider this is because the messages from one speaker are usually not contiguous in a conversation, while DialoGPT can directly produce a response to a message, which is beneficial to BERT on capturing the differences of two messages. (3) During

Iteration	NMI	1-1	Loc ₃	Shen-F	F1
Base	24.96	54.26	60.66	59.16	68.26
1	29.80	56.24	62.39	60.38	68.44
2	29.87	56.33	62.40	60.41	68.47
3	29.71	56.38	62.46	60.44	68.48

Table 2: The performance of the session classifier and the message-pair classifier in each co-training iteration. Columns NMI, 1-1, Loc₃ and Shen-F are for session classifier and Column F1 is for the message-pair classifier. “Base” represents session classifier trained on \mathcal{D}_t and message-pair classifier finetuned on \mathcal{D}_m .

the co-training process, the pseudo pairs retrieved from the predictions of the session classifier can help BERT to achieve a performance close to that of a supervised BERT, which demonstrates the effectiveness of our proposed co-training framework. (4) BERT finetuned with golden message pairs just has a marginal performance advantage compared to the pseudo data \mathcal{D}_m . This is caused by the weakness of two-step methods in which the clustering algorithm is a performance bottleneck (Liu et al., 2020).

In general, end-to-end methods perform much better than two-steps methods as shown in the table, which is in accordance with the conclusions in previous works under supervised settings (Yu and Joty, 2020). The session classifier trained on the pseudo data \mathcal{D}_t can achieve a Shen F score of 59.61, which is +5.29 improvement compared to the supervised BERT in two-step settings. This proves that the model structure and the approach to building \mathcal{D}_t are effective for the task of unsupervised conversation disentanglement. Meanwhile, our proposed co-training framework can further improve the performance of the session classifier and achieve competitive results with the current state-of-the-art supervised method. With further updating during the co-training process, the session classifier raises the NMI score from 24.96 to 29.72 and 1-1 from 54.26 to 56.38. Such a performance gain proves that our co-training framework is an important component in handling unsupervised conversation disentanglement.

Moreover, as we can see in the table, two-step methods have a high MSE on the predicted session numbers, but with the pseudo data \mathcal{D}_m , BERT can achieve performance which is much better than that without finetuning and even comparable with that finetuned on the golden pairs. End-to-end ses-

Type	Hits@1	Hits@2	Hits@5	MRR
None	12.98	23.0	52.81	31.87
Ours	14.50	24.58	54.81	33.29
Gold	17.91	29.97	59.62	37.01

Table 3: The performance on multi-party response selection with disentangled conversations. The first column respective stands for no disentanglement, the disentangled conversations predicted by our method and the golden disentangled conversations.

sion classifier achieves a significant improvement on the MSE by reducing it from 1.4602 to 0.8059, while our proposed co-training framework further improves it to 0.6871, which is close to the performance of the supervised model. It demonstrates that the co-training method can help the session classifier to better understand the semantics in the conversation and thus to more accurately disentangle the conversation into sessions.

5.2.2 Analysis of Co-training

In this section we analyze the iteration process of co-training. Table 2 shows the performance of session classifier in different iterations. We also include in the last column of Table 2 the performance of the message-pair classifier on the task of pair relation prediction.

As we can see, the model performance is improved iteratively as the iteration increases. For the first iteration, the reward r_{m_i} is received from the base message-pair classifier, of which the F1 score on relation prediction is 68.26. After the first iteration, new pseudo pairs will be retrieved from the disentanglement results and used to improve the performance of the message-pair classifier to 68.44. Thus better reward r_{m_i} will be provided to update the session classifier. As shown in the table, with such a co-training procedure, performance of both the session classifier and the message-pair classifier are significantly enhanced.

5.2.3 Performance on Response Selection

Conversation disentanglement is a prerequisite for understanding multi-party conversations. In this section we apply our predicted sessions to the downstream task: multi-party response selection.

We create a response selection dataset based on the Movie Dialogue Dataset. We adopt a LSTM-based network to encode the conversations/sessions, and use attention mechanism to aggregate the information from different sessions as

in Jia et al. (2020). More details of the model and implementation can be found in Appendix A.2.

The results are shown in Table 3. Note that the three experiments are performed on the model with the same number of parameters. We can see that with the disentangled conversations predicted by our method, there is a significant performance gain comparing with the baseline model. Though golden disentanglement can bring the best performance, the annotations are usually expensive to acquire. With our method, a disentanglement scheme can be obtained for better understanding multi-party conversations with no annotation cost.

6 Conclusion

This is the first work to investigate unsupervised conversation disentanglement with deep neural models. We propose a novel approach based on co-training which consists of a message-pair classifier and a session classifier. By iteratively updating the two classifiers with the help of each other, the proposed model attains a performance comparable to that of the state-of-the-art supervised disentanglement methods. Experiments on downstream tasks proves that our method can help better understand multi-party conversations. Our method can be easily adapted to a different assumption, and also it can be extended to other low-resourced scenarios like semi-supervised settings, which we will leave as our future work.

Acknowledgements

We would like to thank the anonymous reviewers for their valuable comments.

References

Avrim Blum and Tom Mitchell. 1998. Combining labeled and unlabeled data with co-training. In *Proceedings of the eleventh annual conference on Computational learning theory*, pages 92–100.

Jun Chen, Chaokun Wang, Heran Lin, Weiping Wang, Zhipeng Cai, and Jianmin Wang. 2017. Learning the structures of online asynchronous conversations. In *International Conference on Database Systems for Advanced Applications*, pages 19–34. Springer.

Muhao Chen, Yingtao Tian, Kai-Wei Chang, Steven Skiena, and Carlo Zaniolo. 2018. Co-training embeddings of knowledge graphs and entity descriptions for cross-lingual entity alignment. In *Proceedings of the Twenty-Seventh International Joint Conference on Artificial Intelligence, IJCAI-18*, pages

3998–4004. International Joint Conferences on Artificial Intelligence Organization.

- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. BERT: Pre-training of deep bidirectional transformers for language understanding. 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 4171–4186, Minneapolis, Minnesota. Association for Computational Linguistics.
- Giacomo Domeniconi, Konstantinos Semertzidis, Vanessa Lopez, Elizabeth M Daly, Spyros Kotoulas, and Gianluca Moro. 2016. A novel method for unsupervised and supervised conversational message thread detection. In *DATA*, pages 43–54.
- 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, Columbus, Ohio. Association for Computational Linguistics.
- Micha Elsner and Eugene Charniak. 2010. Disentangling chat. *Computational Linguistics*, 36(3):389–409.
- 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, Portland, Oregon, USA. Association for Computational Linguistics.
- Shai Erera and David Carmel. 2008. Conversation detection in email systems. In *European Conference on Information Retrieval*, pages 498–505. Springer.
- Gaoyang Guo, Chaokun Wang, Jun Chen, and Pengcheng Ge. 2018. Who is answering to whom? finding “reply-to” relations in group chats with long short-term memory networks. In *Proceedings of the 7th International Conference on Emerging Databases*, pages 161–171. Springer.
- Sepp Hochreiter and Jürgen Schmidhuber. 1997. Long short-term memory. *Neural computation*, 9(8):1735–1780.
- Qi Jia, Yizhu Liu, Siyu Ren, Kenny Zhu, and Haifeng Tang. 2020. Multi-turn response selection using dialogue dependency relations. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 1911–1920, Online. Association for Computational Linguistics.
- 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, New Orleans,

- Louisiana. Association for Computational Linguistics.
- Diederik P. Kingma and Jimmy Ba. 2015. [Adam: A method for stochastic optimization](#). In *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*.
- Jonathan K. Kummerfeld, Sai R. Gouravajhala, Joseph J. Peper, Vignesh Athreya, Chulaka Gunasekara, Jatin Ganhotra, Siva Sankalp Patel, Lazaros C Polymenakos, and Walter Lasecki. 2019. [A large-scale corpus for conversation disentanglement](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 3846–3856, Florence, Italy. Association for Computational Linguistics.
- Hui Liu, Zhan Shi, Jia-Chen Gu, Quan Liu, Si Wei, and Xiaodan Zhu. 2020. [End-to-end transition-based online dialogue disentanglement](#). In *Proceedings of the Twenty-Ninth International Joint Conference on Artificial Intelligence, IJCAI-20*, pages 3868–3874. International Joint Conferences on Artificial Intelligence Organization. Main track.
- Ryan Lowe, Nissan Pow, Iulian Vlad Serban, Laurent Charlin, Chia-Wei Liu, and Joelle Pineau. 2017. Training end-to-end dialogue systems with the ubuntu dialogue corpus. *Dialogue & Discourse*, 8(1):31–65.
- Andrew McCallum and Ben Wellner. 2004. [Conditional models of identity uncertainty with application to noun coreference](#). In *Advances in Neural Information Processing Systems 17 [Neural Information Processing Systems, NIPS 2004, December 13-18, 2004, Vancouver, British Columbia, Canada]*, pages 905–912.
- Kamal Nigam and Rayid Ghani. 2000. Analyzing the effectiveness and applicability of co-training. In *Proceedings of the ninth international conference on Information and knowledge management*, pages 86–93.
- 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, Doha, Qatar. Association for Computational Linguistics.
- Liang Qiu, Yizhou Zhao, Weiyan Shi, Yuan Liang, Feng Shi, Tao Yuan, Zhou Yu, and Song-Chun Zhu. 2020. [Structured attention for unsupervised dialogue structure induction](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 1889–1899, Online. Association for Computational Linguistics.
- Weiyan Shi, Tiancheng Zhao, and Zhou Yu. 2019. [Unsupervised dialog structure learning](#). 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 1797–1807, Minneapolis, Minnesota. Association for Computational Linguistics.
- Richard S Sutton, David A McAllester, Satinder P Singh, Yishay Mansour, et al. 1999. Policy gradient methods for reinforcement learning with function approximation. In *NIPS*, volume 99, pages 1057–1063. Citeseer.
- Ming Tan, Dakuo Wang, Yupeng Gao, Haoyu Wang, Saloni Potdar, Xiaoxiao Guo, Shiyu Chang, and Mo Yu. 2019. [Context-aware conversation thread detection in multi-party chat](#). 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 6456–6461, Hong Kong, China. Association for Computational Linguistics.
- Weishi Wang, Steven C.H. Hoi, and Shafiq Joty. 2020. [Response selection for multi-party conversations with dynamic topic tracking](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 6581–6591, Online. Association for Computational Linguistics.
- Jiawei Wu, Lei Li, and William Yang Wang. 2018. [Reinforced co-training](#). 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 1252–1262, New Orleans, Louisiana. Association for Computational Linguistics.
- Yejun Wu and Douglas W Oard. 2005. Indexing emails and email threads for retrieval. In *Proceedings of the 28th annual international ACM SIGIR conference on Research and development in information retrieval*, pages 665–666.
- Zichao Yang, Diyi Yang, Chris Dyer, Xiaodong He, Alex Smola, and Eduard Hovy. 2016. [Hierarchical attention networks for document classification](#). In *Proceedings of the 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 1480–1489, San Diego, California. Association for Computational Linguistics.
- Tao Yu and Shafiq Joty. 2020. [Online conversation disentanglement with pointer networks](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 6321–6330, Online. Association for Computational Linguistics.
- Ke Zhai and Jason D. Williams. 2014. [Discovering latent structure in task-oriented dialogues](#). In *Proceedings of the 52nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 36–46, Baltimore, Maryland. Association for Computational Linguistics.

Yizhe Zhang, Siqu Sun, Michel Galley, Yen-Chun Chen, Chris Brockett, Xiang Gao, Jianfeng Gao, Jingjing Liu, and Bill Dolan. 2020. [DIALOGPT : Large-scale generative pre-training for conversational response generation](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics: System Demonstrations*, pages 270–278, Online. Association for Computational Linguistics.

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.

A Appendix

A.1 Metric Explanation

We use four metrics in our experiments: Normalized mutual information (NMI), One-to-One Overlap (1-1), Loc₃ and Shen F score (Shen-F). NMI is a normalization of the Mutual Information, which is a method for evaluation of two clusters in the presence of class labels. 1-1 describes how well we can extract whole conversations intact. Loc₃ counts agreements and disagreements within a context window size 3. Shen calculates the F-score for each gold-system conversation pair, finds the max for each gold conversation, and averages weighted by the size of the gold conversation.

A.2 Multi-party Response Selection

Given a conversation $\mathcal{C} = [m_1, \dots, m_N]$ and a candidate message m , the goal of response selection is to decide if message m is correct response to the conversation \mathcal{C} .

We obtain the disentanglement scheme of \mathcal{C} as $\mathbb{T} = \{T^1, T^2, \dots, T^K\}$, where session $T^k = [m_1 \dots, m_{|T^k|}]$. For each session T^k , we encode each message m_i^k within it by a Bidirectional LSTM network and a multilayer perceptron (MLP):

$$\vec{v}_{m_i^k}, \overleftarrow{v}_{m_i^k} = \text{BiLSTM}(m_i^k) \quad (22)$$

$$v_{m_i^k} = \text{MLP}([\vec{v}_{m_i^k}, \overleftarrow{v}_{m_i^k}]) \quad (23)$$

After obtaining sentence embeddings of all the messages in T^k as $[v_{m_1^k}, \dots, v_{m_{|T^k|}^k}]$, we adopt a self attention mechanism (Yang et al., 2016) to calculate the session embedding v_{T^k} by aggregating the information from different messages. Specifically,

$$u_{m_i^k} = \tanh(w \cdot v_{m_i^k} + b) \quad (24)$$

$$\alpha_{m_i^k} = \frac{\exp(u_{m_i^k})}{\sum_i \exp(u_{m_i^k})} \quad (25)$$

$$v_{T^k} = \sum_j \alpha_{m_j^k} v_{m_j^k} \quad (26)$$

where w and b are trainable parameters. In this way we can acquire all the session representation as $\{v_{T^1}, v_{T^2}, \dots, v_{T^K}\}$. Meanwhile, we obtain the candidate message representation as v_m with the same LSTM and MLP in Equation 22-23.

We follow Jia et al. (2020) to aggregate the information from different sessions with the attention

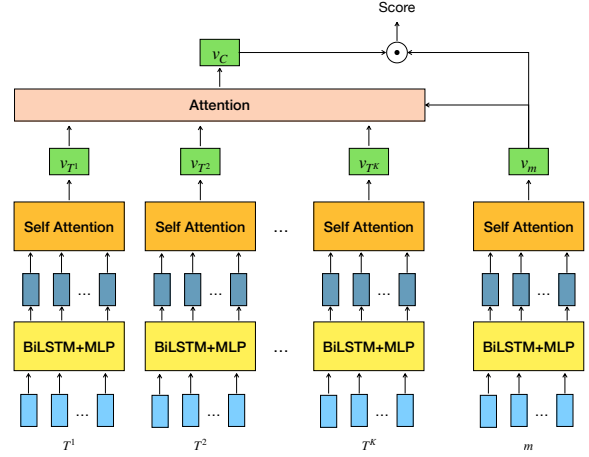


Figure 3: The model structure that incorporating disentangled sessions for the task of response selection.

mechanism:

$$s^k = v_{T^k} \cdot v_m \quad (27)$$

$$w^k = \frac{\exp(s^k)}{\sum_k \exp(s^k)} \quad (28)$$

$$v_C = \sum_k w^k v_{T^k} \quad (29)$$

The final matching score between the conversation and the message is given by:

$$S = v_C \cdot v_m \quad (30)$$

The overall of structure of the method incorporating the disentangled sessions is shown in Figure 3.

For the vanilla model using conversation \mathcal{C} without any disentanglement, we use the LSTM, MLP and self attention as in Equation 22-26 to obtain its vector representation v'_C . And then the matching score is calculated by the dot product between v'_C and v_m .

The whole model is trained to minimize the cross-entropy loss of both positive instances and negative instances.