

Response Selection for Multi-Party Conversations with Dynamic Topic Tracking

Weishi Wang^{§‡}, Steven C.H. Hoi[§], Shafiq Joty^{§‡}

[§]Salesforce Research Asia

[‡]Nanyang Technological University, Singapore

[§]{weishi.wang, shoi, sjoty}@salesforce.com

Abstract

While participants in a multi-party multi-turn conversation simultaneously engage in multiple conversation topics, existing response selection methods are developed mainly focusing on a two-party single-conversation scenario. Hence, the prolongation and transition of conversation topics are ignored by current methods. In this work, we frame response selection as a dynamic topic tracking task to match the topic between the response and relevant conversation context. With this new formulation, we propose a novel multi-task learning framework that supports efficient encoding through large pretrained models with only two utterances at once to perform dynamic topic disentanglement and response selection. We also propose Topic-BERT an essential pretraining step to embed topic information into BERT with self-supervised learning. Experimental results on the DSTC-8 Ubuntu IRC dataset show state-of-the-art results in response selection and topic disentanglement tasks outperforming existing methods by a good margin.¹

1 Introduction

In recent years, with the influx of deep learning methods in natural language processing (NLP), there has been a lot of interests in building effective task-oriented dialogue systems that can assist people in real-world business such as booking tickets, ordering food and solving technical issues (Bui, 2006). *Retrieval-based* response generation that selects a suitable response from a pool of candidates (pre-existing human responses) has become a popular approach to framing dialog. Compared to the *generation-based* systems that generate novel utterances (Serban et al., 2016), retrieval-based systems produce fluent, grammatical and informative responses (Weston et al., 2018; Henderson et al.,

¹Code will be available at <https://github.com/salesforce/TopicBERT>.

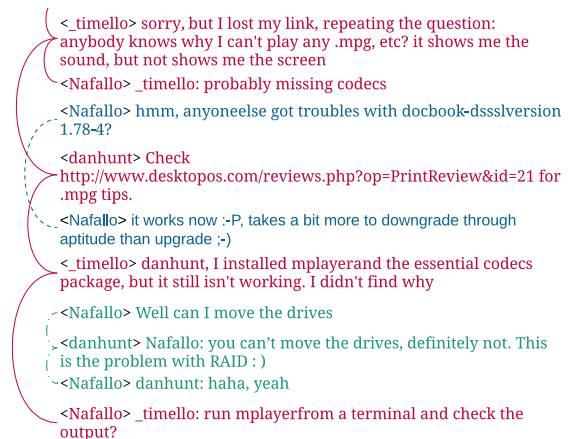


Figure 1: A (truncated) multi-party conversation from Ubuntu IRC log. Curved arrows show the ‘reply-to’ links between utterances. We use different colors to represent different conversation topic clusters.

2019). Also compared to the traditional modular approach, it does not rely on dedicated modules for language understanding, dialog management, and generation, thus simplifying the system design. Due to these reasons, retrieval-based systems have been widely adopted in commercial dialogue systems (Gao et al., 2019; Gunasekara et al., 2019).

Initially, researchers considered response selection in single-turn conversations, where only the last input utterance is considered as the context query (Yan et al., 2016). More recent work deals with multi-turn context, which shows improvements over the single-turn context (Lowe et al., 2015, 2017; Zhou et al., 2016; Chen and Wang, 2019; Gu et al., 2019; Zhou et al., 2018). These methods typically aim to encode the context and the candidate responses in a joint semantic space by capturing short and long range dependencies, and then retrieve the most relevant response by matching the query representation against each candidate’s representation through attentions.

However, most of these works are limited to only

two-party conversations. As dialogue research progresses, it is necessary to study the more generic multi-party multi-turn scenario, which has become very common (e.g., Slack, Whatsapp) with the advent of Internet and mobile devices, and posits a unique set of challenges for the dialog models (Kim et al., 2019; Kummerfeld et al., 2019).

Multiple ongoing conversations seem to occur more naturally in multi-party conversations. For example, consider the conversation excerpt in Figure 1 among three participants, taken from the Ubuntu IRC corpus. There are three ongoing conversation topics as highlighted by different color, and the participants contribute to multiple topics simultaneously (e.g., Nafallo participates in three and danhunt participates in two). An effective response selection method should model such complex conversational topic dynamics in the context, for which existing methods are deficient. In particular, a proper response should match with its context in terms of the same conversation topic, while ignoring other non-relevant topics.

To address the aforementioned challenges in multi-party multi-turn dialog, we frame response selection as a dynamic topic tracking task with the intuition that the topic should remain the same as we go from the context to the response. Our formulation is also supported by the Segmented Discourse Representation Theory (SDRT) of conversations (Asher and Lascarides, 2003). Based on this new formulation, we propose a novel architecture that can incorporate other related dialog tasks such as conversation disentanglement, enabling multi-task learning in a unified framework.

Crucially, our formulation of the task needs to encode only two utterances at a time, thus allowing efficient encoding via large pretrained models like BERT (Devlin et al., 2018). Furthermore, it facilitates pretraining of BERT-like models on topic related sentence pairs to incorporate topic relevance in pretraining, which can be done on large dialog corpora with self-supervised objectives, requiring no manual topic annotations, and can benefit not only response selection but also other dialog tasks. In summary, our contributions are:

- A new formulation of the response selection task with an efficient multi-task learning framework for dynamic topic tracking, which supports efficient encoding with only two utterances at once.
- Incorporate topic prediction and topic disentanglement as auxiliary tasks within the framework.

Based on the similarity of these three tasks, the objective is to match topic (topic prediction) between context utterance and response and track response’s topic (topic disentanglement) across contexts to select an appropriate response.

- Propose Topic-BERT as a pretraining step to embed topic information into BERT, and use a self-supervised approach to generate topic sentence pairs from existing dialogue datasets. The incorporated topic information is shown to be a key step to our topic tracking framework.
- Apply topic attention by using topic embedding as query to obtain utterance-level embeddings for topic prediction. Then self-attention was applied to capture the contextual topic vectors for response selection and topic disentanglement.
- Evaluate the proposed models on the DSTC-8 Ubuntu IRC dataset (Kim et al., 2019), and show state-of-the-art results in both response selection and topic disentanglement outperforming the existing methods by a good margin.

2 Related Work

2.1 Response Selection

A dual encoder framework was proposed to match the context and response (Lowe et al., 2015), and the long short-term memory (LSTM) was utilized to learn the long and short term dependencies among tokens. Beyond tokens, the sentence view matching was introduced by applying a hierarchical recurrent neural network to model sentence level relationships (Zhou et al., 2016). However, context utterances and response are encoded separately without interaction; thus the semantics extracted from context are not based on the response. Recent approaches such as Sequential Matching Network (SMN) (Wu et al., 2019) leverage the contextual information by matching each contextual utterance with response and the multi channel Convolutional Neural Network (CNN) was proposed to generate multiple levels of granularity of matched segment.

These hierarchy-based methods use LSTM to encode the text, which is not cost effective to capture multi-grained segment representation (Lowe et al., 2015; Zhou et al., 2016; Wu et al., 2019). A particular work on sequence-based method stand out in DSTC-7; Enhanced Sequential Inference Model (ESIM) (Chen et al., 2017) achieves the state-of-the-art performance in DSTC-7 by taking advantage of inter-sentence matching (Chen et al.,

2016; Chen and Wang, 2019). It converts multi-turn dialogue setting to natural language inference setting. In addition, transformer-based approach Deep Attention Matching (DAM) solve response selection problem by attention mechanism (Zhou et al., 2018). It utilizes utterance self-attention and context-to-response cross attention to leverage the hidden representation at multi-grained level. Similar to DAM, Multi-hop Selector Network (MSN) was proposed to fuse and select relevant context utterances and match it with the response utterance (Yuan et al., 2019). In addition, Tao et al. (2019) studied the relationship between context utterance and response which indicates that the depth of interaction affect the effectiveness of the model.

Compared to LSTM-based approaches, methods based on transformers (Vaswani et al., 2017) present a promising performance in both accuracy and efficiency (Yang et al., 2020). Devlin et al. (2018) proposed BERT, a transformer-based large-scale pretrained language model, which achieves state-of-the-art performance in different NLP tasks. BERT is also a good match to response selection problem as shown by Vig and Ramea (2019). Our Topic-BERT is initialised with BERT_{base} and post-trained with topic related sentence pairs.

2.2 Hard Context Retrieval

The side effect of multi-speaker multi-turn context is crucial; a lot of noise will be introduced in the context utterances. The speaker and addressee information are essential to decide the structure of conversation, thus can also benefit conversational response selection (Zhang et al., 2017; Le et al., 2019; Hu et al., 2019). A hard context retrieval method was proposed by Wu et al. (2020b) to minimize the context size, while keeping only the utterances whose speaker is the same as the response candidates or referred by the response candidates. However, it cannot guarantee clean context with a single topic of conversation. Indeed, topic tracking is necessary along with hard context retrieval.

2.3 Conversation Disentanglement

Traditional statistical learning based approaches and linguistic features have shown to be effective for conversation disentanglement (Mayfield et al., 2012; Du et al., 2016). Recent methods demonstrate that neural networks could be applied to have a better linguistic representation of the utterances to retrieve relevant conversation. Hand crafted features and pretrained word embeddings are utilized

to predict the link-to relationship between utterances (Kummerfeld et al., 2019). Recently, BERT has been adapted in disentangling task to capture the semantics across utterances (Gu et al., 2020). Also, a masked transformer has been applied to learn the graphical representation of utterances based on the reply-to links (Zhu et al., 2019).

3 Task Formulation

Our Topic-BERT framework combines response selection task with two auxiliary tasks, which are topic prediction and topic disentanglement.

Response Selection Our primary task is response selection in multi-party multi-turn conversations. Let $\mathcal{D}_{rs} = \{(c_i, r_{i,j}, y_{i,j})\}_{i=1}^{|\mathcal{D}_{rs}|}$ is a response selection dataset, where j is the index of a response candidate for a context $c_i = \{u_1, u_2, \dots, u_n\}$ with n utterances. Each utterance $u_i = \{s_i, w_{i,1}, w_{i,2}, \dots, w_{i,m}\}$ starts with its speaker s_i and is composed of m words. Similarly, a response $r_{i,j}$ has a speaker $s_{i,j}$ and composed of n words. $y_{i,j} \in \{0, 1\}$ represents the relevance label. Our goal is to find the relevance ranking score $f_{\theta_r}(c_i, r_{i,j})$ with model parameters θ_r .

Topic Prediction For this (auxiliary) task, we assume a multi-party conversation with a single conversation topic. Let $\mathcal{D}_{tp} = \{(c_i, r_i^+, r_i^-)\}_{i=1}^{|\mathcal{D}_{tp}|}$ is a topic prediction dataset, where r_i^+ is a positive (same conversation) response and r_i^- is a negative (difference conversation) response for context c_i . For our training purposes, each utterance pair from the same context constitutes (c_i, r_i^+) , whereas an utterance pair from different contexts constitutes (c_i, r_i^-) . Our goal is to train a binary classifier $g_{\theta_t}(c_i, r_i) \in \{0, 1\}$ with model parameters θ_t .

Topic Disentanglement In this (auxiliary) task, our goal is to disentangle single conversations from a multi-party conversation based on topics. For a given conversation context $c_i = \{u_1, u_2, \dots, u_n\}$, a set of pairwise “reply-to” annotations $\mathcal{R} = \{(u_c, u_p)_1, \dots, (u_c, u_p)_{|\mathcal{R}|}\}$ is given, where u_p is a parent of child u_c . Our task is to compute a reply-to score $h_{\theta_d}(u_i, u_j)$ for $j \leq i$ that indicates the score for u_j being the parent of u_i , with model parameters θ_d . The individual conversations can then be constructed by following the reply-to links. Note that an utterance u_i can point to itself, which we call *self-link*. Self-links are either start of a conversation or a system message, and they play a crucial

role in identifying the conversation clusters.

4 Our Topic-BERT Framework

Our framework for response selection aims to track how the conversation topics change from one utterance to another and use it for ranking the candidate responses. As shown in Fig. 2, we encode an utterance u_k from the context $c_i = \{u_1, u_2, \dots, u_n\}$ along with a candidate response $r_{i,j}$ using our pretrained Topic-BERT encoder (§4.1). The contextual token representations in Topic-BERT encode topic relevance between the tokens of u_k and the tokens of $r_{i,j}$, while the [CLS] representation captures utterance-level topic relevance. We use the [CLS] representation as query to attend over the token representations to further enforce topic relevance in the attended topic vector t_k .

We repeat this encoding process for the n utterances in the context c_i to get n different topic vectors $T_j = \{t_1, \dots, t_n\}$ that model $r_{i,j}$'s topic relevance to the each of context utterances. These topic representations are then used for the prediction tasks – topic prediction, disentanglement, and response selection. Response selection is our main task, while the other two tasks are auxiliary and optional. Since our Topic-BERT encodes two utterances at a time, the encoding process is efficient and can be used to encode larger context. The core component of our framework is the Topic-BERT pretraining as we describe next.

4.1 Topic-BERT Pretraining

One crucial advantage of our topic-based task formulation is that it allows us to pretrain BERT directly on a very relevant task in a self-supervised way, without requiring any human annotation. In other words, our goal is to pretrain BERT such that it can be used to encode relevant topic information for our task(s). For this, we assume that a single-threaded conversation between two or more participants covers a single topic and the utterance pairs in that thread can be used to pretrain our Topic-BERT with relevant self-supervised objectives.

To collect such single-threaded conversational data in an opportunistic way, we can simply adopt the heuristics (unsupervised) used by Lowe et al. (2015) to collect the popular Ubuntu Dialogue Corpus from multi-threaded chatlogs. Alternatively, we can extract two-party conversations from other sources as done in previous work (Henderson et al., 2019; Wu et al., 2020a). In our experiments, we use

the data from DSTC-8 task 1 (Kim et al., 2019), which was automatically collected from Ubuntu chat logs. This dataset contains detached speaker-visible conversations between two or more participants from the Ubuntu IRC channel.

To pretrain Topic-BERT, we first initialise it with the pretrained uncased BERT_{base} (Devlin et al., 2018). We treat the training setting similar to our *topic prediction* task in §3. Formally, the pretraining dataset is $\mathcal{D}_{pr} = \{(u_i, r_i^+, r_{i,j}^-)\}_{i=1}^{|\mathcal{D}_{pr}|}$, where each utterance pair from the same conversation (including the true response) constitutes a positive pair (u_i, r_i^+) , and for each such positive pair we randomly sample 4 negative responses $(r_{i,j}^-)$ from the 100 candidate pool to balance the positive and negative ratio. We (re)train Topic-BERT on \mathcal{D}_{pr} with two self-supervised objectives as follows.

Masked Language Modeling (MLM) We follow the same MLM training of the original BERT (Devlin et al., 2018) by masking 15% of the input tokens at random, and replacing the masked word with [MASK] token at 80% of the time, with a random word at 10% of the time, and with the original word at 10% of the time. The MLM objective is only applied to the positive samples.

Same Topic Prediction (STP) Each training pair $((u_i, r_i^+) \text{ or } (u_i, r_{i,j}^-))$ is fed into the Topic-BERT as $([\text{CLS}], [u_1], [\text{SEP}], [u_2], [\text{SEP}])$. Similar to the original BERT's Next Sentence Prediction (NSP) task, the position embedding, segment embedding and token embedding are added together to get input layer token representations. The token representations are then passed through multiple transformer (Vaswani et al., 2017) encoder layers, where each layer is comprised of a self-attention and a feed-forward sublayer. Different from the original BERT, Topic-BERT uses the [CLS] representation to predict whether the training instance is a positive (same topic) pair or a negative (different topic) pair. Thus, the [CLS] representation encodes topic relationship between the two utterances and will be used as the topic-aware contextual embedding to determine whether the two utterances are matched in topic.

4.2 Topic-BERT Multi-Task Framework

As shown in Fig. 2(b), the encoded representations from our Topic-BERT are passed through a topic attention layer (§4.2.1) to get the corresponding topic vectors, which are then used for the end tasks.

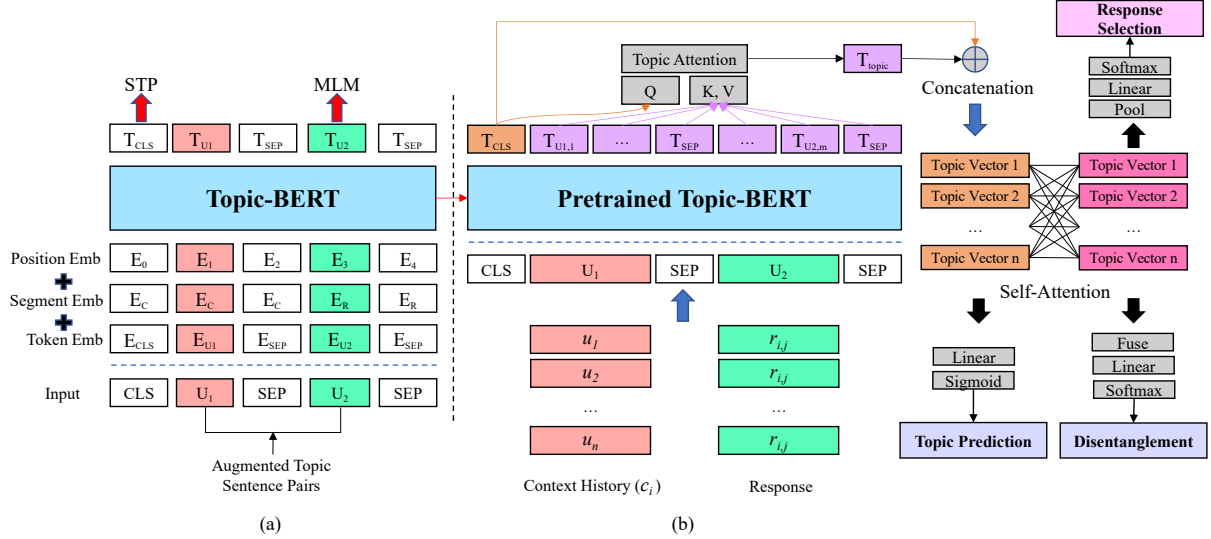


Figure 2: Overview of Topic-BERT architecture. (a) Topic-BERT pretraining with topic sentence pairs to incorporate utterance-utterance topic relationship. (b) Our multi-task framework which uses the pretrained Topic-BERT to enhance topic information in the encoded representations to support three downstream tasks – response selection as the main task while topic prediction and disentanglement as two auxiliary (optional) tasks.

4.2.1 Topic Attention Layer

We apply an attention layer to enhance topic information in the encoded vector. We use the Topic-BERT’s [CLS] representation T_{CLS} as query to attend to the remaining K tokens $\{T_j\}_{j=1}^K$:

$$e_j = v_a^T \tanh(\mathbf{W}_a T_{CLS} + \mathbf{U}_a T_j); \quad (1)$$

$$a_j = \frac{\exp(e_j)}{\sum_{j=1}^K \exp(e_i)} \quad (2)$$

$$T_{\text{topic}} = \sum_{j=1}^K a_j T_j \quad (3)$$

where v_a , \mathbf{W}_a and \mathbf{U}_a are trainable parameters. The concatenation of T_{topic} and T_{CLS} constitutes the final topic vector, *i.e.*, $\mathbf{t} = [T_{CLS}; T_{\text{topic}}]$. We repeat this encoding process for the n utterances in the context $c_i = \{u_1, u_2, \dots, u_n\}$ by pairing each with the candidate response $r_{i,j}$ to get n different topic vectors $\mathbf{T}_j = \{\mathbf{t}_1, \dots, \mathbf{t}_n\}$. \mathbf{T}_j represents $r_{i,j}$ ’s topic relevance to the context utterances, which will be fed to the task-specific layers.

4.2.2 Topic Prediction

Topic prediction is done for each utterance-response pair $(u_k, r_{i,j})$ for all $u_k \in c_i$ to decide whether u_k and $r_{i,j}$ should be in the same topic (§3). The Topic-BERT encoded topic vector corresponding to the $(u_k, r_{i,j})$ pair is $\mathbf{t}_k \in \mathbf{T}_j$. We define the binary topic classification model as:

$$g_{\theta_t}(u_k, r_{i,j}) = \text{sigmoid}(\mathbf{w}_p^T \mathbf{t}_k) \quad (4)$$

where \mathbf{w}_p is the task-specific parameter. We use a binary cross entropy loss computed as:

$$\mathcal{L}_{\text{topic}} = -y \log(g_{\theta_t}) - (1 - y) \log(1 - g_{\theta_t}) \quad (5)$$

where $y \in \{0, 1\}$ is the ground truth indicating same or different topic. Note that topic prediction is an auxiliary task intended to help our main task of response selection, as we describe next.

4.2.3 Response Selection

In response selection, our goal is to measure relevance of a candidate response $r_{i,j}$ with respect to the context c_i . For this, we first apply the same hard context retrieval method proposed by Wu et al. (2020b) to filter out irrelevant utterances and to reduce the context size. Then, we put each context utterance paired with the response $r_{i,j}$ as the input to Topic-BERT to compute the corresponding topic vectors \mathbf{T}_j through the topic attention layer.

We pass the topic vectors $\mathbf{T}_j \in \mathbb{R}^{n \times d}$ through a scaled dot-product self-attention layer (Vaswani et al., 2017) to learn all-pair topic relevance at the utterance level. Formally,

$$\mathbf{T}'_j = \text{softmax}\left(\frac{(\mathbf{T}_j \mathbf{W}_q)(\mathbf{T}_j \mathbf{W}_k)^\top}{\sqrt{d}}\right)(\mathbf{T}_j \mathbf{W}_v) \quad (6)$$

where $\{\mathbf{W}_q, \mathbf{W}_k, \mathbf{W}_v\} \in \mathbb{R}^{n \times d}$ are the query, key and value parameters, respectively, and d denotes the hidden dimension of 768.

We add a max-pooling layer to select the most important information followed by a linear layer

and a softmax to compute the relevance score of the response $r_{i,j}$ with the context c_i . Formally,

$$f_{\theta_r}(u_k, r_{i,j}) = \text{softmax}(\mathbf{W}_r(\text{maxpool}(\mathbf{T}'_j))) \quad (7)$$

where \mathbf{W}_r is the task-specific parameter. We use the standard cross entropy loss defined as:

$$\mathcal{L}_{\text{rs}} = - \sum_{i,j} \mathbb{1}(y_{i,j}) \log(f_{\theta_r}) \quad (8)$$

where $\mathbb{1}(y_{i,j})$ is the one-hot encoding of the ground truth label.

4.2.4 Topic Disentanglement

For topic disentanglement (§3), our goal is to find the “reply-to” links between the utterances (including the candidate response) to track which utterance is replying to which previous utterance.

For training on topic disentanglement, we simulate a sliding window over the entire (entangled) conversation. Each window constitutes a context $c_i = \{u_1, u_2, \dots, u_n\}$ and the model is trained to find the parent of u_n in c_i , in other words, we try to find the reply-to link (u_n, u_{n_p}) for $1 \leq n_p \leq n$.

For the input to our Topic-BERT (Fig. 2b), we treat u_n as the response, thus allowing also response-response (u_n, u_n) interactions through Topic-BERT’s encoding layers to facilitate *self-link* predictions (the fact that u_n can point to itself).

In the task-specific layer for disentanglement, we take the self-attended topic vectors $\mathbf{T}'_j = \{\mathbf{t}'_1, \dots, \mathbf{t}'_n\}$ as input, and separate it into two parts: context topic vectors encapsulated in $\mathbf{T}'_c = \{\mathbf{t}'_1, \dots, \mathbf{t}'_{n-1}\} \in \mathbb{R}^{(n-1) \times d}$ and the response topic vector $\mathbf{t}'_n \in \mathbb{R}^d$. In order to model high-order interactions between the response and context utterances, we compute the differences and element-wise products between them (Chen and Wang, 2019). We duplicate the response message \mathbf{t}'_n to obtain $\mathbf{T}'_r \in \mathbb{R}^{(n-1) \times d}$ and concatenate them as:

$$\mathbf{T}'' = [\mathbf{T}'_r, \mathbf{T}'_c, \mathbf{T}'_r \odot \mathbf{T}'_c, \mathbf{T}'_r - \mathbf{T}'_c] \quad (9)$$

Then, we compute the *reply-to* distribution as: $h_{\theta_d}(u_n, c_i) = \text{softmax}(\mathbf{T}'' \mathbf{w}_d) \in \mathbb{R}^{n \times 1}$, and optimize with the following cross-entropy loss:

$$\mathcal{L}_{\text{dis}} = - \sum_{j=1}^n \mathbb{1}(y_j) \log(h_{\theta_d}) \quad (10)$$

For inference, we compute $\arg \max_j h_{\theta_d}(u_n, c_j)$.

Tasks@Track2	Train Set	Val Set	Test Set
Task 1 (# Dialog)	225,367	4,827	5,529
Task 2 (# Dialog)	112,262	9,565	9,027
Task 4 (# Link)	69,395	2,607	5,187

Table 1: DSTC-8 Ubuntu Dataset Statistics.

4.2.5 Multi-task Learning

We jointly train the three tasks: response selection, topic prediction and topic disentanglement, which share the same topic attention weights to benefit each other. Response selection should benefit from dynamic topic prediction and disentanglement. Similarly, topic prediction and disentanglement should benefit from the response prediction. The overall loss is a combination of the three task losses from Equations 5, 8, and 10:

$$\mathcal{L} = \alpha \mathcal{L}_{\text{rs}} + \beta \mathcal{L}_{\text{topic}} + \gamma \mathcal{L}_{\text{dis}} \quad (11)$$

where α , β , and γ are parameters which are chosen from $[0, 0.1, 0.2, \dots, 1]$ by optimizing our model response selection accuracy on dev dataset.

5 Experiments

In this section, we present our experiments, including the datasets, experimental setup, evaluation metrics, and the results with analysis.

Datasets and Setup Considering multi-party conversations, we adopt a publicly available Ubuntu dataset from DSTC-8 track 2 “NOESIS II: Predicting Responses”(Kim et al., 2019). This dataset consists of four tasks and we use the datasets from three of them, including Task 1: single-topic multi-party dialogues for response selection; Task 2: a long Ubuntu chat log with multi-party conversations of multiple on-going topics simultaneously, which is ideal for our main response selection evaluation; Task 4: multi-party chat with link annotations (used for disentanglement task). Table 1 shows the dataset statistics. More details about the datasets, experimental setups and training details can be found in Appendix.

Evaluation Metrics DSTC-8 Track 2 considered a range of metrics for comparing models. We follow their evaluation metrics and the details could be found in Appendix.

5.1 Experiment I: Response Selection

Baseline Models. We compare the proposed Topic-BERT approach with several existing and state-of-the-art approaches for response selection:

- **BERT.** We adopt the vanilla pretrained uncased BERT_{base}² as the base model, and follow (Gu et al., 2020) to post-train BERT_{base} for 10 epochs on DSTC-Task 1 (response selection in a single-topic dialog). We take the whole context with the response as one input sequence. We then finetune it on Task 2’s response selection for 10 more epochs. More details can be found in Appendix.
- **ToD-BERT.** This is a domain-specific pretrained BERT from Wu et al. (2020a), which is pre-trained on a combination of 9 Task-oriented Dialogue datasets and surpasses BERT in several downstream response selection tasks.
- **BERT-ESIM.** This model ensembles both ESIM (Chen et al., 2017) and BERT with gradient boosting classifier, and ranks the second best in DSTC-8 response selection (Dario Bertero, 2020).
- **Adapt-BERT.** This is based on BERT model with task-related pretraining and context modeling through hard and soft context modeling, and ranks as top-1 in the DSTC-8 response selection challenge (Wu et al., 2020b).

Results. From Table 2, we can see that our Topic-BERT model outperforms the baselines by a large margin. By examining our model in detail, we found that our context filtering, self-supervised topic training and topic attention contribute positively to our model, boosting the metric of Recall@1 from 0.287 (BERT_{base}) to 0.696 (Topic-BERT with standalone response selection task). This shows our topic pretraining with task related data improves BERT for response selection task.

Furthermore, the performance continues to increase from 0.696 to 0.710, when we jointly train response selection and topic prediction (2nd last row), validating an effective utilization of topic information in selecting response. Then we replace topic prediction with disentanglement, which further improves from 0.710 to 0.720, showing response selection can utilize topic tracing by sharing the connection of utterances. Finally, our Topic-BERT with the multi-task learning achieves the best result (0.726) and significantly outperform the prior state-of-the-art Adapt-BERT in DSTC-8 response selection task (Kim et al., 2019).

We further compute BLEU4 SacreBLEU (Post, 2018) for the incorrectly selected responses by Topic-BERT and ToD-BERT. From Table 3, we

²<https://github.com/huggingface/transformers>

Model	Recall@1	Recall@5	Recall@10	MRR
BERT _{base}	0.287	0.503	0.572	0.351
BERT _{+post-train}	0.532	0.797	0.840	0.677
ToD-BERT	0.588	0.823	0.885	0.691
Adapt-BERT	0.706	0.916	0.957	0.799
Topic-BERT	0.726	0.930	0.970	0.807
-TP	0.720	0.927	0.964	0.803
-D	0.710	0.924	0.960	0.800
-TP -D	0.696	0.910	0.950	0.790

Table 2: Response selection results on DSTC-8 Ubuntu. “-TP” means our model excluding topic prediction loss and “-D” means excluding topic disentanglement loss. Adapt-BERT results are obtained from (Wu et al., 2020b), other DSTC-8 released baselines are in Appendix.

Model	BLEU4	Precision@N-gram			
		N = 1	N = 2	N = 3	N = 4
ToD-BERT	0.67	7.568	1.894	0.218	0.065
Topic-BERT	0.75	7.876	2.032	0.250	0.078

Table 3: BLEU4 and N-gram precision are calculated using SacreBLEU on incorrectly selected responses.

see that responses retrieved by Topic-BERT are more relevant even if they are not the top one.

5.2 Experiment II: Topic Prediction

This experiment aims to examine how significant our Topic-BERT can improve over the baselines on the topic prediction task, which is important for both response selection and topic disentanglement.

Baseline Models.

- **BERT.** We use our post-trained BERT_{base} from §5.1 and fine-tune it on Task 1 topic sentence pairs as our BERT baseline for topic prediction.
- **ToD-BERT.** We adopt our post-trained ToD-BERT and fine-tune it with our obtained topic sentences pairs as the ToD-BERT baseline.

Results. Table 4 gives the topic prediction results on DSTC-8 task-1. From the results, we can see that our Topic-BERT outperforms the baselines BERT and ToD-BERT significantly in the topic prediction task. Compared with our pretrained Topic-BERT without fine-tuning (last row), the proposed topic attention further enhances the topic matching of two utterances by improving the F-score by 1.5% (from 0.813 to 0.828). Joint training with response selection or disentanglement tasks show similar effect on topic prediction tasks, and the contextual topic information sharing by Topic-BERT

Model	Precision	Recall	F-Score
BERT	0.523	0.482	0.502
ToD-BERT	0.626	0.563	0.593
Topic-BERT	0.890	0.847	0.868
-D	0.891	0.845	0.867
-RS	0.889	0.840	0.864
-D -RS	0.866	0.793	0.828
w/o FT	0.848	0.781	0.813

Table 4: Topic prediction results on DSTC-8 Ubuntu. “w/o FT” means our Topic-BERT without fine-tuning, “-RS” means our model excluding the Response Selection loss, “-D” means excluding Disentanglement loss.

Model	Precision	Recall	F-Score
BERT	0.431	0.417	0.424
MH BERT	0.539	0.517	0.528
ToD-BERT	0.612	0.603	0.607
Feed-Forward	0.748	0.718	0.733
Topic-BERT	0.754	0.725	0.739
-TP	0.749	0.727	0.737
-RS	0.705	0.692	0.698
-TP -RS	0.689	0.678	0.683

Table 5: Disentanglement results on DSTC-8 Ubuntu. “-RS” means our model excluding Response Selection loss, and “-TP” means excluding Topic Prediction loss.

multi-task model add a marginal improvement in topic prediction. Compared with vanilla BERT, ToD-BERT (Wu et al., 2020a) makes substantial improvement for the topic prediction task, but not as significant as ours. This further confirms the importance and efficacy of our learning scheme. Meanwhile, if we compare our pretrained Topic-BERT without fine-tuning (last row) with the BERT model that does not use STP (first row), the significant improvement gives us an impression on how much our Topic-BERT benefits from the STP loss.

5.3 Experiment III: Disentanglement

This experiment aims to examine how well can Topic-BERT tackle the topic disentanglement task.

Baseline Models.

- **BERT & ToD-BERT** We use our fine-tuned BERT and ToD-BERT models in §5.2 as our baselines by taking the history of utterances $(u_1, \dots, u_{n-1}, u_n)$ and pair each with the current utterance u_n itself from a dialogue as input. Following (Gu et al., 2020), A single-layer BiLSTM is applied to extract the cross message semantics of [CLS] outputs. Then we take the differences and element-wise products (Eq. 9) between the

Model	Recall10@1	Recall10@2	Recall10@5
DL2R	0.626	0.783	0.944
Multi View	0.662	0.801	0.951
SMN _{dynamic}	0.726	0.847	0.961
AK-DE-biGRU	0.747	0.868	0.972
DUA	0.752	0.868	0.962
DAM	0.767	0.874	0.969
IMN	0.777	0.888	0.974
ESIM	0.796	0.894	0.975
MRFN _{FLS}	0.786	0.886	0.976
BERT _{base}	0.817	0.904	0.977
BERT-DPT	0.851	0.924	0.984
Topic-BERT	0.861	0.933	0.985

Table 6: Response selection results on Ubuntu Corpus v1. All other results are from (Whang et al., 2019).

history and current utterance. Finally, a feedforward layer is used for link prediction.

- **Feed-Forward.** This is the baseline model³ from DSTC-8 task organizers that has the best result for task 4 (Kummerfeld et al., 2019), which is trained by employing a two-layer feed-forward neural network on a set of 77 *hand engineered features* combined with word average embeddings from pretrained Glove embeddings.
- **Masked Hierarchical (MH) BERT.** This is a two-stage BERT proposed by Zhu et al. (2019) to model the conversation structure, in which the low-level BERT is to capture the utterance-level contextual representation between utterances, and the high-level BERT is to model the conversation structure with an ancestor masking approach to avoid irrelevant connections.

Results. From the results in Table 5, we can see that our Topic-BERT achieves the best result and outperforms all the BERT based baselines significantly. This shows our multi-task learning can enrich the link relationship for improving disentanglement together with topic prediction and response selection. The improvement of Topic-BERT over the baseline model using feed-forward network and hand-crafted features is relatively less, but our approach is able to avoid manual feature engineering. Many of these features are dataset/domain specific and they do not generalize across datasets/domains.

5.4 Experiment IV: Evaluation on New Task

Finally, we examine our Topic-BERT’s transferability on a new task based on another Ubuntu Corpus

³<https://github.com/dstc8-track2/NOESIS-II/tree/master/subtask4>

v1 dataset by comparing with various state-of-the-art response selection methods in Table 6. Ubuntu Corpus V1 contains 1M train set, 500K validation and 500K test set (Lowe et al., 2015).

Baseline Models. Here we mainly introduce the state-of-the-art baseline: BERT-DPT (Whang et al., 2019), which fine-tunes BERT by optimizing the domain post-training (DPT) loss comprising both NSP and MLM objectives for response selection. Details of other baselines can be found in Appendix.

Results. Our Topic-BERT with standalone response selection task fine-tuned on Ubuntu Corpus v1 outperforms the state-of-the-art BERT-DPT, improved by about 1% for Recall₁₀@1. This result shows that the learned topic relevance in Topic-BERT can be potentially transferable to a novel task, the topic information influences the response selection positively, and our utterance-level topic tracking is effective for response selection.

6 Conclusion

This paper presented a new formulation of response selection in multi-party conversations from a novel dynamic topic tracking perspective. Based on our new formulation, we propose Topic-BERT for response selection in multi-party conversations, which consists of two steps: (1) a topic-based pre-training to embed topic information into BERT with self-supervised learning, and (2) a multi-task learning on our pretrained model by jointly training response selection and dynamic topic prediction and disentanglement tasks. Empirically the proposed Topic-BERT achieved the state-of-the-art results on the DSTC8 Ubuntu IRC datasets.

References

N. Asher and A. Lascarides. 2003. *Logics of Conversation*. Cambridge University Press.

Trung Bui. 2006. Multimodal dialogue management-state of the art. *Surface Science - SURFACE SCI*.

Qian Chen and Wen Wang. 2019. Sequential attention-based network for noetic end-to-end response selection. *ArXiv*, abs/1901.02609.

Qian Chen, Xiaodan Zhu, Zhen-Hua Ling, Si Wei, and Hui Jiang. 2016. *Enhancing and combining sequential and tree LSTM for natural language inference*. *CoRR*, abs/1609.06038.

Qian Chen, Xiaodan Zhu, Zhen-Hua Ling, Si Wei, Hui Jiang, and Diana Inkpen. 2017. Enhanced lstm for natural language inference. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1657–1668.

Kenichi Yokote Makoto Iwayama Kenji Nagamatsu Dario Bertero, Takeshi Homma. 2020. *Model ensembling of esim and bert for dialogue response selection*.

Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2018. *BERT: pre-training of deep bidirectional transformers for language understanding*. *CoRR*, abs/1810.04805.

Wenchao Du, Pascal Poupart, and Wei Xu. 2016. Discovering conversational dependencies between messages in dialogs.

Jianfeng Gao, Michel Galley, and Lihong Li. 2019. *Neural approaches to conversational ai*. *Foundations and Trends® in Information Retrieval*, 13(2-3):127–298.

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*.

Jia-Chen Gu, Zhen-Hua Ling, and Quan Liu. 2019. *Interactive matching network for multi-turn response selection in retrieval-based chatbots*. In *Proceedings of the 28th ACM International Conference on Information and Knowledge Management, CIKM '19*, page 2321–2324, New York, NY, USA. Association for Computing Machinery.

Chulaka Gunasekara, Jonathan K. Kummerfeld, Lazaros Polymenakos, and Walter Lasecki. 2019. *DSTC7 task 1: Noetic end-to-end response selection*. In *Proceedings of the First Workshop on NLP for Conversational AI*, pages 60–67, Florence, Italy. Association for Computational Linguistics.

Matthew Henderson, Iñigo Casanueva, Nikola Mrkšić, Pei-Hao Su, Tsung-Hsien Wen, and Ivan Vulić. 2019. *Convert: Efficient and accurate conversational representations from transformers*.

Wenpeng Hu, Zhangming Chan, Bing Liu, Dongyan Zhao, Jinwen Ma, and Rui Yan. 2019. Gsn: A graph-structured network for multi-party dialogues. *arXiv preprint arXiv:1905.13637*.

Seokhwan Kim, Michel Galley, Chulaka Gunasekara, Sungjin Lee, Adam Atkinson, Baolin Peng, Hannes Schulz, Jianfeng Gao, Jinchao Li, Mahmoud Adada, et al. 2019. The eighth dialog system technology challenge. *arXiv preprint arXiv:1911.06394*.

Jonathan K. Kummerfeld, Sai R. Gouravajhala, Joseph Peper, Vignesh Athreya, Chulaka Gunasekara, Jatin Ganhotra, Siva Sankalp Patel, Lazaros Polymenakos,

- and Walter S. Lasecki. 2019. A large-scale corpus for conversation disentanglement. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*.
- Ran Le, Wenpeng Hu, Mingyue Shang, Zhenjun You, Lidong Bing, Dongyan Zhao, and Rui Yan. 2019. Who is speaking to whom? learning to identify utterance addressee in multi-party conversations. 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 1909–1919.
- Ryan Lowe, Nissan Pow, Iulian Serban, and Joelle Pineau. 2015. [The Ubuntu dialogue corpus: A large dataset for research in unstructured multi-turn dialogue systems](#). In *Proceedings of the 16th Annual Meeting of the Special Interest Group on Discourse and Dialogue*, pages 285–294, Prague, Czech Republic. Association for Computational Linguistics.
- Ryan Thomas Lowe, Nissan Pow, Iulian Serban, Laurent Charlin, Chia-Wei Liu, and Joelle Pineau. 2017. Training end-to-end dialogue systems with the ubuntu dialogue corpus. *DD*, 8:31–65.
- Elijah Mayfield, David Adamson, and Carolyn Penstein Rosé. 2012. [Hierarchical conversation structure prediction in multi-party chat](#). In *Proceedings of the 13th Annual Meeting of the Special Interest Group on Discourse and Dialogue*, pages 60–69, Seoul, South Korea. Association for Computational Linguistics.
- Matt Post. 2018. [A call for clarity in reporting BLEU scores](#). In *Proceedings of the Third Conference on Machine Translation: Research Papers*, pages 186–191, Belgium, Brussels. Association for Computational Linguistics.
- Iulian Vlad Serban, Ryan Lowe, Laurent Charlin, and Joelle Pineau. 2016. [Generative deep neural networks for dialogue: A short review](#).
- Chongyang Tao, Wei Wu, Can Xu, Wenpeng Hu, Dongyan Zhao, and Rui Yan. 2019. One time of interaction may not be enough: Go deep with an interaction-over-interaction network for response selection in dialogues. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 1–11.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin. 2017. [Attention is all you need](#). *CoRR*, abs/1706.03762.
- Jesse Vig and Kalai Ramea. 2019. Comparison of transfer-learning approaches for response selection in multi-turn conversations.
- Jason Weston, Emily Dinan, and Alexander Miller. 2018. [Retrieve and refine: Improved sequence generation models for dialogue](#). In *Proceedings of the 2018 EMNLP Workshop SCAI: The 2nd International Workshop on Search-Oriented Conversational AI*, pages 87–92, Brussels, Belgium. Association for Computational Linguistics.
- Taesun Whang, Dongyub Lee, Chanhee Lee, Kisu Yang, Dongsuk Oh, and HeuiSeok Lim. 2019. Domain adaptive training bert for response selection. *arXiv preprint arXiv:1908.04812*.
- Chien-Sheng Wu, Steven Hoi, Richard Socher, and Caiming Xiong. 2020a. [Tod-bert: Pre-trained natural language understanding for task-oriented dialogues](#). *arXiv preprint arXiv:2004.06871*.
- Shuangzhi Wu, Yufan Jiang, Xu Wang, Wei Miao, Zhenyu Zhao, Xie Jun, and Mu Li. 2020b. [Enhancing response selection with advanced context modeling and post-training](#).
- Yu Wu, Wei Wu, Chen Xing, Can Xu, Zhoujun Li, and Ming Zhou. 2019. [A sequential matching framework for multi-turn response selection in retrieval-based chatbots](#). *Computational Linguistics*, 45(1):163–197.
- Rui Yan, Yiping Song, and Hua Wu. 2016. [Learning to respond with deep neural networks for retrieval-based human-computer conversation system](#). In *Proceedings of the 39th International ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR '16*, page 55–64, New York, NY, USA. Association for Computing Machinery.
- Yiying Yang, Kaijie Zhou, Xuan Li, Dawei Zhu, Xiaoyuan Yao, and Jianping Shen. 2020. [Transformer-based semantic matching model for noetic response selection](#).
- Chunyan Yuan, Wei Zhou, Mingming Li, Shangwen Lv, Fuqing Zhu, Jizhong Han, and Songlin Hu. 2019. Multi-hop selector network for multi-turn response selection in retrieval-based chatbots. 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 111–120.
- Rui Zhang, Honglak Lee, Lazaros Polymenakos, and Dragomir Radev. 2017. Addressee and response selection in multi-party conversations with speaker interaction rnns. *arXiv preprint arXiv:1709.04005*.
- Xiangyang Zhou, Daxiang Dong, Hua Wu, Shiqi Zhao, Dianhai Yu, Hao Tian, Xuan Liu, and Rui Yan. 2016. [Multi-view response selection for human-computer conversation](#). In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 372–381, Austin, Texas. Association for Computational Linguistics.
- Xiangyang Zhou, Lu Li, Daxiang Dong, Yi Liu, Ying Chen, Wayne Xin Zhao, Dianhai Yu, and Hua Wu. 2018. [Multi-turn response selection for chatbots with deep attention matching network](#). In *Proceedings of the 56th Annual Meeting of the Association*

for Computational Linguistics (Volume 1: Long Papers), pages 1118–1127, Melbourne, Australia. Association for Computational Linguistics.

Henghui Zhu, Feng Nan, Zhiguo Wang, Ramesh Nallapati, and Bing Xiang. 2019. Who did they respond to? conversation structure modeling using masked hierarchical transformer. *arXiv preprint arXiv:1911.10666*.