

HETERMPC: A Heterogeneous Graph Neural Network for Response Generation in Multi-Party Conversations

Jia-Chen Gu^{1*}†, Chao-Hong Tan¹†, Chongyang Tao², Zhen-Hua Ling¹,
Huang Hu², Xiubo Geng², Daxin Jiang²‡

¹National Engineering Research Center for Speech and Language Information Processing,
University of Science and Technology of China, Hefei, China

²Microsoft, Beijing, China

{gujc, chtan}@mail.ustc.edu.cn, zhling@ustc.edu.cn,
{chotao, huahu, xigeng, djiang}@microsoft.com

Abstract

Recently, various response generation models for two-party conversations have achieved impressive improvements, but less effort has been paid to multi-party conversations (MPCs) which are more practical and complicated. Compared with a two-party conversation where a dialogue context is a sequence of utterances, building a response generation model for MPCs is more challenging, since there exist complicated context structures and the generated responses heavily rely on both interlocutors (i.e., speaker and addressee) and history utterances. To address these challenges, we present HeterMPC, a heterogeneous graph-based neural network for response generation in MPCs which models the semantics of utterances and interlocutors simultaneously with two types of nodes in a graph. Besides, we also design six types of meta relations with node-edge-type-dependent parameters to characterize the heterogeneous interactions within the graph. Through multi-hop updating, HeterMPC can adequately utilize the structural knowledge of conversations for response generation. Experimental results on the Ubuntu Internet Relay Chat (IRC) channel benchmark show that HeterMPC outperforms various baseline models for response generation in MPCs.

1 Introduction

Enabling dialogue systems to converse naturally with humans is a challenging yet intriguing problem of artificial intelligence and has attracted increasing attention due to its promising potentials and alluring commercial values (Kepuska and Bohouta, 2018; Berdasco et al., 2019; Zhou et al., 2020). A large number of researchers have focused on building dialogue generation models with various neural networks. At first, researchers mostly




Figure 1: Illustration of a graphical information flow in an MPC. Pink rectangles denote utterances and blue circles denote interlocutors. Each solid line represents the “replied-by” relationship between two utterances. Each dashed line indicates the speaker of an utterance.

focused on dialogues between two participants (Shang et al., 2015; Serban et al., 2016; Wen et al., 2017; Young et al., 2018). Recently, researchers have paid more attention to a more practical and challenging scenario involving more than two participants, which is well known as multi-party conversations (MPCs) (Ouchi and Tsuboi, 2016; Zhang et al., 2018; Le et al., 2019; Hu et al., 2019b; Wang et al., 2020b; Gu et al., 2021). Utterances in a two-party conversation are posted one by one between two interlocutors, constituting a *sequential* information flow. Different from that, utterances in an MPC can be spoken by anyone and address anyone else in this conversation, which constitutes a *graphical* information flow as shown in Figure 1.

Although sequence-to-sequence (Seq2Seq) models (Sutskever et al., 2014; Serban et al., 2016) are effective at modeling sequential dialogues, they fall short of modeling graph-structured ones. To overcome this drawback, Hu et al. (2019b) first proposed a graph-structured network (GSN) to encode utterances based on the graph topology rather than the sequence of their appearances. The graph established in GSN was homogeneous, where nodes represented only utterances. However, interlocutors are also important components of MPCs. There exist complicated interactions between interlocutors, and between an utterance

*Work done during the internship at Microsoft.

†Equal contribution.

‡Corresponding author.

and an interlocutor. Furthermore, when passing messages over a graph, a bidirectional information flow algorithm was designed for GSN. Since both the forward and backward information flows employed the same model structure and parameters, this algorithm cannot distinguish the “reply” or “replied-by” relations between two connected utterance nodes. Also, information flows along both directions are independently propagated, so that a graph node cannot be jointly updated at a single propagation step.

On account of above issues, we propose a heterogeneous graph-based neural network for response generation in MPCs, named HeterMPC. First, a heterogeneous graph is designed which employs two types of nodes to represent utterances and interlocutors respectively. Different from previous methods that built a homogeneous graph modeling only utterances, utterances and interlocutors are modeled simultaneously in HeterMPC, so that the complicated interactions between interlocutors, between utterances, and between an interlocutor and an utterance can be explicitly described. In order to characterize the heterogeneous attention over each (*source*, *edge*, *target*) triple, model parameters dependent on both types of nodes and edges are introduced when calculating attention weights and passing messages. Specifically, we introduce six types of meta relations for modeling different edges including “reply” and “replied-by” between two utterances, “speak” and “spoken-by” between an utterance and a speaker, and “address” and “addressed-by” between an utterance and an addressee. With these node-edge-type-dependent structures and parameters, HeterMPC can better utilize the structural knowledge of conversations for node representation and response generation than conventional homogeneous graphs. Finally, Transformer is employed as the backbone of HeterMPC and its model parameters can be initialized with PLMs to take advantage of the recent breakthrough on pre-training.

We evaluate HeterMPC on the Ubuntu Internet Relay Chat (IRC) channel benchmark released by Hu et al. (2019b). Experimental results show that HeterMPC outperforms GSN (Hu et al., 2019b), GPT-2 (Radford et al., 2019), BERT (Devlin et al., 2019) and BART (Lewis et al., 2020) by significant margins in terms of both automated and human evaluation metrics, achieving a new state-of-the-art performance for response generation in MPCs.

In summary, our contributions in this paper are three-fold: 1) To the best of our knowledge, this paper is the first exploration of using heterogeneous graphs for modeling conversations; 2) A Transformer-based heterogeneous graph architecture is introduced for response generation in MPCs, in which two types of nodes, six types of meta relations, and node-edge-type-dependent parameters are employed to characterize the heterogeneous properties of MPCs; 3) Experimental results show that our proposed model achieves a new state-of-the-art performance of response generation in MPCs on the Ubuntu IRC benchmark.

2 Related Work

Multi-Party Conversation Existing methods on building dialogue systems can be generally categorized into generation-based (Shang et al., 2015; Serban et al., 2016; Wen et al., 2017; Young et al., 2018; Zhang et al., 2020) or retrieval-based approaches (Lowe et al., 2015; Wu et al., 2017; Zhou et al., 2018; Tao et al., 2019a,b; Gu et al., 2019, 2020). In this paper, we study the task of response generation in MPCs, where in addition to utterances, interlocutors are also important components who play the roles of speakers or addressees. Previous methods have explored retrieval-based approaches for MPCs. For example, Ouchi and Tsuboi (2016) proposed the dynamic model which updated speaker embeddings with conversation streams. Zhang et al. (2018) proposed speaker interaction RNN which updated speaker embeddings role-sensitively. Wang et al. (2020b) proposed to track the dynamic topic in a conversation. Gu et al. (2021) proposed jointly learning “who says what to whom” in a unified framework by designing self-supervised tasks during pre-training. On the other hand, Hu et al. (2019b) explored generation-based approaches by proposing a graph-structured network, the core of which was an utterance-level graph-structured encoder.

Heterogeneous Graph Neural Network Early studies on graph neural networks (GNNs) focused on homogeneous graphs where a whole graph is composed of a single type of nodes. However, graphs in real-world applications usually come with multiple types of nodes, namely heterogeneous information networks (HINs) or heterogeneous graphs (Sun and Han, 2012). Recently, researchers have attempted to extend GNNs to model heterogeneity. For example, Zhang et al.

(2019) adopted different RNNs for different types of nodes to integrate multi-modal features. Wang et al. (2019) extended graph attention networks by maintaining different weights for different meta-path-defined edges. Hu et al. (2020) proposed heterogeneous graph Transformer (HGT) to model heterogeneity by maintaining dedicated representations for different types of nodes and edges. In addition, heterogeneous graphs have also been applied to many NLP tasks, such as multi-hop reading comprehension (Tu et al., 2019), text classification (Hu et al., 2019a) and document summarization (Wang et al., 2020a).

Previous studies have verified the superiority of modeling MPCs with homogeneous graphs considering only utterances. We claim that it is indeed necessary to model a complex information flow in MPCs shown in Figure 1 with a heterogeneous graph, since a homogeneous one cannot explicitly model the relationships of multiple utterances spoken by or addressing an interlocutor. Nowadays, HINs have been widely used in many NLP tasks. To the best of our knowledge, this paper makes the first attempt to build a heterogeneous graph-based neural network considering utterances and interlocutors simultaneously for response generation in MPCs. In addition, we introduce many task-specific modelings for MPCs such as graph construction and node updating which will be elaborated in the model section.

3 Problem Formulation

The task of response generation in MPCs is to generate an appropriate response \bar{r} given the conversation history, the speaker of a response, and which utterance the response is going to reply to, which can be formulated as:

$$\begin{aligned} \bar{r} &= \operatorname{argmax}_r \log P(r|\mathbb{G}) \\ &= \operatorname{argmax}_r \sum_{k=1}^{|r|} \log P(r_k|\mathbb{G}r_{<k}). \end{aligned} \quad (1)$$

Here, \mathbb{G} is a heterogeneous graph containing both history conversation and the response to be generated. The speaker and addressee of the response are known and its contents are masked. The response tokens are generated in an autoregressive way. r_k and $r_{<k}$ stand for the k -th token and the first $(k-1)$ tokens of response r respectively. $|r|$ is the length of r .




Figure 2: Illustration of the two types of nodes and six types of edges in a heterogeneous conversation graph. This example demonstrates that I_j speaks U_n replying U_m that is spoken-by I_i .

We will introduce how to construct the graph and how to model the probability in Eq. (1) given the built graph in the next section.

4 HeterMPC Model

HeterMPC adopts an encoder-decoder architecture consisting of stacked encoder and decoder layers for graph-to-sequence learning (Yao et al., 2020). The graph encoder is designed to capture conversation structures and output the representations of all nodes in a graph that are fed to the decoder for response generation.

4.1 Graph Construction

A heterogeneous graph is constructed to explicitly model the complicated interactions between interlocutors, between utterances, and between an interlocutor and an utterance in an MPC. This graph models utterances and interlocutors by considering them as two types of nodes under a unified framework. Given an MPC instance composed of M utterances and I interlocutors, a heterogeneous graph $\mathbb{G}(\mathbb{V}, \mathbb{E})$ is constructed. Specifically, \mathbb{V} is a set of $M+I$ nodes. Each node denotes either an utterance or an interlocutor. $\mathbb{E} = \{e_{p,q}\}_{p,q=1}^{M+I}$ is a set of directed edges. Each edge $e_{p,q}$ describes the connection from node p to node q .

Inspired by Sun et al. (2011, 2012), we introduce six types of meta relations $\{\text{reply}, \text{replied-by}, \text{speak}, \text{spoken-by}, \text{address}, \text{addressed-by}\}$ to describe the directed edge between two graph nodes as illustrated in Figure 2. Specifically, if an utterance represented by node n replies another utterance represented by node m , the edge $e_{n,m} = \text{reply}$ and the reversed edge $e_{m,n} = \text{replied-by}$. If an utterance represented by node m is spoken by an interlocutor represented by node i , $e_{i,m} = \text{speak}$ and $e_{m,i} = \text{spoken-by}$. If an utterance represented by node n addresses an interlocutor represented by

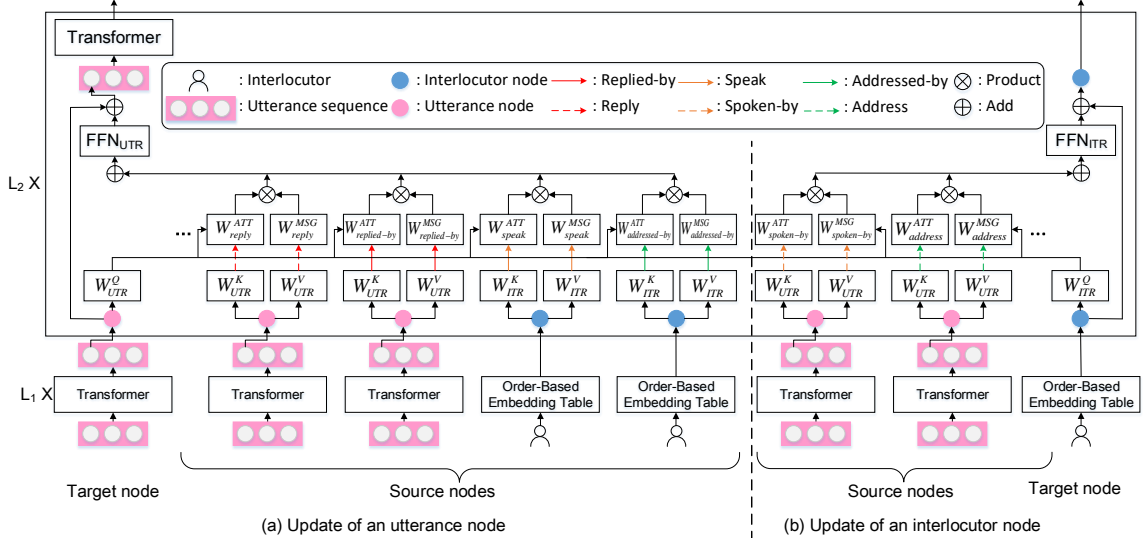


Figure 3: Model architecture of HeterMPC for (a) update of an utterance node and (b) update of an interlocutor node. “UTR” and “ITR” are abbreviations of “utterance” and “interlocutor” respectively. The set of W_*^* denotes the node-edge-type-dependent parameters.

node i , $e_{n,i} = \textit{address}$ and $e_{i,n} = \textit{addressed-by}$. In other cases, $e_{p,q} = \textit{NULL}$ indicating that there is no connection between these two nodes. Note that it is necessary to distinguish the bidirectional edges between every two nodes that indicate the active and passive tense information respectively.

4.2 Node Initialization

In HeterMPC, each node is represented as a vector. These vectors are first initialized individually without considering graph edges.

Utterances When encoding utterances, a [CLS] token is inserted at the start of each utterance, denoting the utterance-level representation. Besides, a [SEP] token is also inserted at the end of each utterance (Devlin et al., 2019). Then each utterance is encoded individually by stacked Transformer encoder layers through the self-attention mechanism to derive the contextualized utterance representations.¹ The output of a Transformer encoder layer is used as the input of the next layer. Readers can refer to Vaswani et al. (2017) for details of Transformer. Formally, the calculation for an utterance at the l -th Transformer layer is denoted as:

$$\mathbf{H}_m^{l+1} = \text{TransformerEncoder}(\mathbf{H}_m^l), \quad (2)$$

¹In our experiments, BERT or BART was selected to initialize the utterance encoder layers of HeterMPC. Then, the built HeterMPC models were compared with the baseline models directly finetuning BERT or BART, respectively. It is worth noting that the utterance encoder layers of HeterMPC can also be initialized by other types of PLMs, and the comparison across PLMs is not the focus of this paper.

where $m \in \{1, \dots, M\}$, $l \in \{0, \dots, L_1 - 1\}$, L_1 denotes the number of Transformer layers for initialization, $\mathbf{H}_m^l \in \mathbb{R}^{k_m \times d}$, k_m denotes the length of an utterance and d denotes the dimension of embedding vectors.

Interlocutors Different from an utterance composed of a sequence of tokens, an interlocutor is directly represented with an embedding vector. Interlocutors in a conversation are indexed according to their speaking order and the embedding vector for each interlocutor is derived by looking up an order-based interlocutor embedding table (Gu et al., 2020) that is updated during end-to-end learning. The first interlocutors in all conversation sessions share the same embedding vector in the interlocutor embedding table, so do all the second interlocutors.² Thus, this order-based embedding table can be shared across the training, validation and testing sets, and there is no need to estimate an embedding vector for each specific interlocutor in the dataset.

4.3 Node Updating

As shown in Figure 3, the initialized node representations are updated by feeding them into the built graph for absorbing context information (Kipf and Welling, 2017; Velickovic et al., 2018; Yun

²In our experiments, the maximum interlocutor number was set to 10 and an embedding table sized 10*768 was learned during training. We did study initializing the embedding vector of an interlocutor node by averaging the representations of all utterance nodes it speaks, but no further improvement can be achieved.

et al., 2019). We calculate heterogeneous attention weights between connected nodes and pass messages over the graph in a node-edge-type-dependent manner, inspired by introducing parameters to maximize feature distribution differences for modeling heterogeneity (Schlichtkrull et al., 2018; Wang et al., 2019; Zhang et al., 2019; Hu et al., 2020). After collecting the information from all source nodes to a target node, a node-type-dependent feed-forward network (FFN) followed by a residual connection (He et al., 2016) is employed to aggregate the information. Then, in order to let each token in an utterance have access to the information from other utterances, an additional Transformer layer is placed for utterance nodes specifically. L_2 denotes the number of iterations for updating both utterance and interlocutor nodes.

4.3.1 Heterogeneous Attention

Since the representations of two types of nodes are initialized in different ways, node-type-dependent linear transformations are first applied to node representations before attention calculation so that the two types of nodes share similar feature distributions (Wang et al., 2019; Hu et al., 2020). Meanwhile, each of the six relation types is assigned a separate linear projection so that the semantic relationship between two connected nodes can be accurately described when calculating attention weights. The forward and backward information flows between them can also be distinguished.

Formally, let the triple (s, e, t) denote an edge e connecting a source node s to a target node t . The representations of the source and target nodes at the l -th iteration³ are denoted as \mathbf{h}_s^l and \mathbf{h}_t^l , serving as a *key* (K) vector and a *query* (Q) vector of attention calculation respectively. Then, the heterogeneous attention weight $w^l(s, e, t)$ before normalization for this triple is calculated as:

$$\mathbf{k}^l(s) = \mathbf{h}_s^l \mathbf{W}_{\tau(s)}^K + \mathbf{b}_{\tau(s)}^K, \quad (3)$$

$$\mathbf{q}^l(t) = \mathbf{h}_t^l \mathbf{W}_{\tau(t)}^Q + \mathbf{b}_{\tau(t)}^Q, \quad (4)$$

$$w^l(s, e, t) = \mathbf{k}^l(s) \mathbf{W}_{e,s,t}^{ATT} \mathbf{q}^l(t)^T \frac{\mu_{e,s,t}}{\sqrt{d}}. \quad (5)$$

Here, $\tau(s), \tau(t) \in \{UTR, ITR\}$ distinguish utterance (*UTR*) and interlocutor (*ITR*) nodes. Eqs. (3) and (4) are node-type-dependent linear transformations. Eq. (5) contains an edge-type-dependent linear projection $\mathbf{W}_{e,s,t}^{ATT}$ where $\mu_{e,s,t}$ is an adaptive

³For an utterance, the representation for the [CLS] token is extracted as the utterance-level representation.

factor scaling to the attention. All $\mathbf{W}^* \in \mathbb{R}^{d \times d}$ and $\mathbf{b}^* \in \mathbb{R}^d$ are parameters to be learnt.

4.3.2 Heterogeneous Message Passing

When passing the message of a source node that serves as a *value* (V) vector to a target node, node-edge-type-dependent parameters are also introduced considering the heterogeneous properties of nodes and edges. Mathematically:

$$\mathbf{v}^l(s) = \mathbf{h}_s^l \mathbf{W}_{\tau(s)}^V + \mathbf{b}_{\tau(s)}^V, \quad (6)$$

$$\bar{\mathbf{v}}^l(s) = \mathbf{v}^l(s) \mathbf{W}_{e,s,t}^{MSG}, \quad (7)$$

where $\bar{\mathbf{v}}^l(s)$ is the passed message and all $\mathbf{W}^* \in \mathbb{R}^{d \times d}$ and $\mathbf{b}^* \in \mathbb{R}^d$ are parameters to be learnt.

4.3.3 Heterogeneous Aggregation

For a target node, the messages passed from all its connected source nodes need to be aggregated. A softmax function is applied to normalize the attention weights and then the messages from all source nodes are summarized as:

$$\bar{\mathbf{h}}_t^l = \sum_{s \in S(t)} \text{softmax}(w^l(s, e, t)) \bar{\mathbf{v}}^l(s), \quad (8)$$

where $S(t)$ denotes the set of source nodes for the target node t . Then the summarized message $\bar{\mathbf{h}}_t^l$ is aggregated with the original node representation \mathbf{h}_t^l using a node-type-dependent FFN followed by a residual connection (He et al., 2016) as:

$$\mathbf{h}_t^{l+1} = \text{FFN}_{\tau(t)}(\bar{\mathbf{h}}_t^l) + \mathbf{h}_t^l, \quad (9)$$

where the output \mathbf{h}_t^{l+1} is used as the input of the next iteration of node updating. One iteration can be viewed as a single-step information propagation along edges. When stacking L_2 iterations, a node can attend to other nodes up to L_2 hops away.

A specific consideration on utterance nodes is that the tokens except [CLS] in an utterance have no access to other utterances during the node updating process introduced above. To overcome this disadvantage and derive more contextualized utterance representations, an additional Transformer layer (Vaswani et al., 2017) is further placed for utterance nodes as shown in Figure 3. In detail, at the l -th iteration, the representations of an utterance node before and after node updating, i.e., \mathbf{h}_t^l and \mathbf{h}_t^{l+1} , are concatenated and then compressed by a linear transformation as:

$$\hat{\mathbf{h}}_t^{l+1} = [\mathbf{h}_t^l; \mathbf{h}_t^{l+1}] \mathbf{W}_{com} + \mathbf{b}_{com}, \quad (10)$$

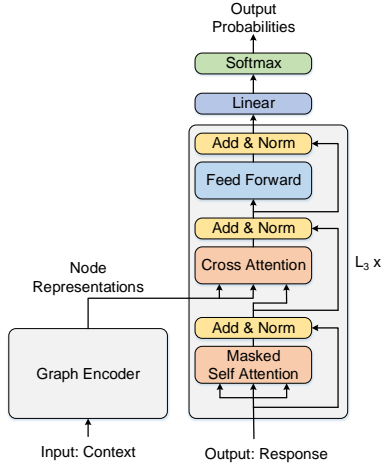


Figure 4: The decoder architecture of HeterMPC.

where $W_{com} \in \mathbb{R}^{2d \times d}$ and $b_{com} \in \mathbb{R}^d$ are parameters. Then, \hat{h}_t^{l+1} replaces the representation of [CLS] (i.e., h_t^l) in the sequence representations of the whole utterance. Finally, the updated sequence representations are fed into the additional Transformer layer for another round of intra-utterance self-attention, so that the context information learnt by the [CLS] representation can be shared with other tokens in the utterance.

4.4 Decoder

The decoder is composed of a stack of identical layers as shown in Figure 4. We follow the standard implementation of Transformer decoder to generate responses. In each decoder layer, a masked self-attention operation is first performed where each token cannot attend to future tokens to avoid information leakage. Furthermore, a cross-attention operation over the node representations of the graph encoder output is performed to incorporate graph information for decoding. It is notable that a residual connection along with layer normalization is followed by each attention operation.

5 Experiments

5.1 Datasets

We evaluated our proposed method on the Ubuntu IRC benchmark used in Hu et al. (2019b). The data processing script provided by Hu et al. (2019b) was employed to derive the dataset.⁴ In this dataset, both speaker and addressee labels were included for each utterance in a session. When testing, the

⁴We contacted the authors of Hu et al. (2019b) to obtain the data processing script. As they claimed, it was an updated version which was a little different from that used in their paper. Thus, we re-implemented all baselines on this updated dataset to ensure fair comparison.

speaker and addressee information was both given for response generation, i.e., the system knew who would speak next and which utterance should be responded to following the graph structure. It contained 311,725/5,000/5,000 dialogues in the training/validation/testing sets respectively.

5.2 Baseline Models

We compared our proposed methods with as many MPC models as possible. Considering that there are only a few research papers in this field, several recent advanced models were also adapted to provide sufficient comparisons. Finally, we compared with the following baseline models: **(1) RNN-based Seq2Seq** (Sutskever et al., 2014) took all utterances except the target utterance to generate as input, which were sorted according to their posting time and concatenated. Thus, structured conversations were converted into sequential ones. Seq2Seq modeling with attention was performed as that in Sutskever et al. (2014); Bahdanau et al. (2015) on the concatenated utterances. **(2) Transformer** (Vaswani et al., 2017) took the same input utterances as those used for the Seq2Seq model. **(3) GPT-2** (Radford et al., 2019) was a uni-directional pre-trained language model. Following its original concatenation operation, all context utterances and the response were concatenated with a special [SEP] token as input for encoding. **(4) BERT** (Devlin et al., 2019) concatenated all context utterances and the response similarly as those for GPT-2. To adapt BERT for response generation, a special masking mechanism was designed to avoid response information leakage during encoding. Concretely, each token in the context utterances attended to all tokens in the context utterances, while each token in the response cannot attend to future tokens in the utterance. **(5) GSN** (Hu et al., 2019b) achieved the state-of-the-art performance on MPCs. The core of GSN was an utterance-level graph-structured encoder. **(6) BART** (Lewis et al., 2020) was a denoising autoencoder using a standard Transformer-based architecture, trained by corrupting text with an arbitrary noising function and learning to reconstruct the original text. In our experiments, a concatenated context started with <s> and separated with </s> were fed into the encoder, and a response were fed into the decoder.

5.3 Evaluation Metrics

To ensure all experimental results were comparable, we used the same automated and human evaluation

Metrics Models	BLEU-1	BLEU-2	BLEU-3	BLEU-4	METEOR	ROUGE _L
Seq2Seq (LSTM) (Sutskever et al., 2014)	7.71	2.46	1.12	0.64	3.33	8.68
Transformer (Vaswani et al., 2017)	7.89	2.75	1.34	0.74	3.81	9.20
GSN (Hu et al., 2019b)	10.23	3.57	1.70	0.97	4.10	9.91
GPT-2 (Radford et al., 2019)	10.37	3.60	1.66	0.93	4.01	9.53
BERT (Devlin et al., 2019)	10.90	3.85	1.69	0.89	4.18	9.80
HeterMPC _{BERT}	12.61	4.55	2.25	1.41	4.79	11.20
HeterMPC _{BERT} w/o. node types	11.76	4.09	1.87	1.12	4.50	10.73
HeterMPC _{BERT} w/o. edge types	12.02	4.27	2.10	1.30	4.74	10.92
HeterMPC _{BERT} w/o. node and edge types	11.60	3.98	1.90	1.18	4.20	10.63
HeterMPC _{BERT} w/o. interlocutor nodes	11.80	3.96	1.75	1.00	4.31	10.53
BART (Lewis et al., 2020)	11.25	4.02	1.78	0.95	4.46	9.90
HeterMPC _{BART}	12.26	4.80	2.42	1.49	4.94	11.20
HeterMPC _{BART} w/o. node types	11.22	4.06	1.87	1.04	4.57	10.63
HeterMPC _{BART} w/o. edge types	11.52	4.27	2.05	1.24	4.78	10.90
HeterMPC _{BART} w/o. node and edge types	10.90	3.90	1.79	1.01	4.52	10.79
HeterMPC _{BART} w/o. interlocutor nodes	11.68	4.24	1.91	1.03	4.79	10.65

Table 1: Performance of HeterMPC and ablations on the test set in terms of automated evaluation. Numbers in bold denote that the improvement over the best performing baseline is statistically significant (t-test with p -value < 0.05).

Metrics Models	Score	Kappa
Human	2.81	0.55
GSN (Hu et al., 2019b)	2.00	0.50
BERT (Devlin et al., 2019)	2.19	0.42
BART (Lewis et al., 2020)	2.24	0.44
HeterMPC _{BERT}	2.39	0.50
HeterMPC _{BART}	2.41	0.45

Table 2: Human evaluation results of HeterMPC and some selected systems on a randomly sampled test set.

metrics as those used in previous work (Hu et al., 2019b). Hu et al. (2019b) used the evaluation package released by Chen et al. (2015) including BLEU-1 to BLEU-4, METEOR and ROUGE_L, which was also used in this paper.⁵ Human evaluation was conducted to measure the quality of the generated responses in terms of three independent aspects: 1) relevance, 2) fluency and 3) informativeness. Each judge was asked to give three binary scores for a response, which were further summed up to derive the final score ranging from 0 to 3.

5.4 Training Details

Model parameters were initialized with pre-trained weights of *bert-base-uncased* released by Wolf et al. (2020). The AdamW method (Loshchilov

⁵<https://github.com/tylin/coco-caption>

and Hutter, 2019) was employed for optimization. The learning rate was initialized as $6.25e-5$ and was decayed linearly down to 0. The max gradient norm was clipped down to 1.0. The batch size was set to 16 with 8 gradient accumulation steps. The maximum utterance length was set to 50. The number of layers for initializing utterance representations L_1 was set to 9, and the number of layers for heterogeneous graph iteration L_2 was set to 3. L_1 and L_2 were validated on the validation set. The number of decoder layers L_3 was set to 6, achieving the best performance out of {2, 4, 6, 8} on the validation set. The strategy of greedy search was performed for decoding. The maximum length of responses for generation was also set to 50. All experiments were run on a single GeForce RTX 2080 Ti GPU. The maximum number of epochs was set to 15, taking about 40 hours. The validation set was used to select the best model for testing. All code was implemented in the PyTorch framework⁶ and are published to help replicate our results.⁷

5.5 Evaluation Results

In our experiments, BERT and BART were selected to initialize HeterMPC. HeterMPC_{BERT} denoted that the utterance encoder was initialized with BERT and the decoder was randomly initialized. HeterMPC_{BART} denoted the encoder and decoder

⁶<https://pytorch.org/>

⁷<https://github.com/lxchtan/HeterMPC>

were initialized by those of BART respectively.

Automated Evaluation Table 1 presents the evaluation results of HeterMPC_{BERT}, HeterMPC_{BART} and previous methods on the test set. Each model ran four times with identical architectures and different random initializations, and the best out of them was reported. We ran the code released by Hu et al. (2019b) to reproduce the results of GSN for a fair comparison.⁸ The results show that both HeterMPC_{BERT} and HeterMPC_{BART} outperformed all baselines in terms of all metrics. HeterMPC_{BERT} outperformed GSN by 2.38% BLEU-1 and 0.44% BLEU-4, and outperformed GPT-2 by 2.24% BLEU-1 and 0.48% BLEU-4. HeterMPC_{BART} outperformed GSN by 2.03% BLEU-1 and 0.52% BLEU-4, and outperformed GPT-2 by 1.89% BLEU-1 and 0.56% BLEU-4. Furthermore, HeterMPC_{BERT} outperformed BERT by 1.71% BLEU-1 and 0.52% BLEU-4, and HeterMPC_{BART} outperformed BART by 1.01% BLEU-1 and 0.54% BLEU-4, illustrating the importance of modeling MPC structures.

To further verify the effectiveness of our proposed methods, ablation tests were conducted as shown in Table 1. First, all nodes or edges were considered equivalently by employing the same linear transformations in Eqs. (3) to (9) for all node or edge types without distinguishing them. The drop in performance illustrates the effectiveness of the node-edge-type-dependent parameters. On the other hand, interlocutor nodes were removed out of a graph and only the meta relations of *reply* and *replied-by* were left. The drop in performance illustrates the importance of modeling interactions between utterances and interlocutors, and the effectiveness of the heterogeneous architecture.

Human Evaluation Table 2 presents the human evaluation results on a randomly sampled test set. 200 samples were evaluated and the order of evaluation systems were shuffled. Three graduate students were asked to score from 0 to 3 (3 for the best) and the average scores were reported. The Fleiss’s kappa value (Fleiss, 1971) for each model was also reported, indicating the inter-judge moderate agreement during evaluation. It can be seen that HeterMPC_{BERT} and HeterMPC_{BART} achieved higher subjective quality scores than the baselines. Their kappa values were also higher than the BERT and BART baselines, respectively.

⁸<https://github.com/morning-dews/GSN-Dialogues>




Figure 5: Performance of HeterMPC_{BERT} under different numbers of iterations (L_2) on the test set.




Figure 6: Performance of BERT and HeterMPC_{BERT} on test samples with different session lengths.

5.6 Analysis

The impact of numbers of iterations (L_2). Figure 5 illustrates how the performance of HeterMPC_{BERT} changed with respect to different numbers of iterations (L_2) on the test set. It can be seen that the performance of HeterMPC_{BERT} was significantly improved as L_2 increased at the beginning, which shows the effectiveness of incorporating the contextual information between nodes. Then, the performance was stable and dropped slightly. The reason might be that models begin to overfit due to a larger set of parameters.

The impact of conversation length.⁹ Figure 6 illustrates how the performance of HeterMPC_{BERT} changed according to the test samples with different session lengths. As the session length increased, the performance of HeterMPC_{BERT} dropped less than that of BERT, showing superiority of our method on dealing with longer conversations.

⁹This experiment can also be considered as exploring the impact of interlocutor number, since more context utterances indicate that more interlocutors are involved implicitly.

Speaker	Utterance	Addressee
Case 1		
I.1	geev: in gparted now?	-
I.2	there is no such command in my computer	I.1
I.1	open a terminal, type: sudo apt-get install gparted	I.2
I.1	after: gparted aksks for your pw, opens a gui	I.2
I.2	ok if found fail i will contact you	I.1
I.3	does it say "to install type sudo apt-get install gparted"? (Human)	I.2
	do you have remote access to remote desktop? (GSN)	
	ok, i will try that. (BERT)	
	you can use the command "sudoapt-get install gparted" to install the gparted livecd. (HeterMPC_{BERT})	
	i'm not sure what you mean by "in gparted" (BART)	
	you need to install gparted, it is a gui tool (HeterMPC_{BART})	
Case 2		
I.1	got a bit of a mount problem. fstab is still the same but the secondary harddrive shows as busy when doing mount	-
I.2	you can see what is mounted with : mount	I.1
I.1	yes it is not mounted	I.2
I.3	how would i check what module the cam uses?	I.2
I.2	lsmod less will show you te loaded modules	I.3
I.1	lmount - t ext3 filepath fails cause already mounted or busy	I.2
I.2	then try to mount it at command line (Human)	I.1
	i'm not sure how to do that (GSN)	
	i'm not sure what the problem is (BERT)	
	you need to mount it as a mount point (HeterMPC_{BERT})	
	i'm not sure what the problem is (BART)	
	you need to check the filepath file (HeterMPC_{BART})	

Table 3: The response generation results of two test samples. "I." is the abbreviation of "interlocutor". We kept original texts without manual corrections.

Case Study. Case studies were conducted by randomly sampling two MPC instances as shown in Table 3. Given the conversation graph of the first case, the response to generate addresses I.2. Thus, the information relevant to I.2 should be collected. We can see that "gparted" in the first utterance is two hops away from I.2 (the first utterance is replied by the second utterance which is spoken by I.2), and this word in the fourth utterance and "install gparted" in the third utterance are both one hop away from I.2 (these

two utterances directly address I.2). The responses generated by HeterMPC_{BERT} and HeterMPC_{BART} both contain these keywords, showing that it can capture the conversation graph information accurately and generate a human-like response. However, due to the lack of the interlocutor information and the conversation structure, GSN generated an irrelevant response. BERT generated a response which seems replying to the third utterance. Although BART captured "gparted", it failed to handle the action "install". In the second case, we can see that the responses generated by GSN, BERT and BART are general and useless while HeterMPC_{BERT} and HeterMPC_{BART} can still generate a suitable response. Due to the complicated interactions between utterances and interlocutors, the conversation flow might be led by some unnecessary information, which shows the importance of making models aware of the conversation structure.

Robustness. The addressee labels are important for constructing a graph used in HeterMPC. This kind of label is commonly available in real life such as "A@B" labels in group chatting, Twitter, Reddit and various forums that denote speaker A talking to addressee B. However, addressee labels of a part of utterances are missing in the existing MPC datasets since a speaker may forget to specify an addressee. HeterMPC is robust since utterances without addressee labels can be assigned with a general addressee label "To all interlocutors". We leave evaluation on other datasets in future work.

6 Conclusion

We present HeterMPC to model complicated interactions between utterances and interlocutors in MPCs with a heterogeneous graph. Two types of graph nodes and six types of edges are designed. Node-edge-type-dependent parameters are introduced for better utilizing the structural knowledge of conversations during node updating. Results show that HeterMPC outperforms baselines by significant margins, achieving a new state-of-the-art performance for response generation in MPCs on the Ubuntu IRC benchmark. In the future, we will explore better ways of maximizing feature distribution differences to model heterogeneity.

Acknowledgements

We thank anonymous reviewers for their valuable comments.

References

- Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio. 2015. [Neural machine translation by jointly learning to align and translate](#). In *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*.
- Ana Berdasco, Gustavo López, Ignacio Díaz-Oreiro, Luis Quesada, and Luis A. Guerrero. 2019. [User experience comparison of intelligent personal assistants: Alexa, google assistant, siri and cortana](#). In *13th International Conference on Ubiquitous Computing and Ambient Intelligence, UCAmI 2019, Toledo, Spain, December 2-5, 2019*, volume 31 of *MDPI Proceedings*, page 51. MDPI.
- Xinlei Chen, Hao Fang, Tsung-Yi Lin, Ramakrishna Vedantam, Saurabh Gupta, Piotr Dollár, and C. Lawrence Zitnick. 2015. [Microsoft COCO captions: Data collection and evaluation server](#). *CoRR*, abs/1504.00325.
- 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, NAACL-HLT 2019, Minneapolis, MN, USA, June 2-7, 2019, Volume 1 (Long and Short Papers)*, pages 4171–4186.
- Joseph L Fleiss. 1971. Measuring nominal scale agreement among many raters. *Psychological bulletin*, 76(5):378–382.
- Jia-Chen Gu, Tianda Li, Quan Liu, Zhen-Hua Ling, Zhiming Su, Si Wei, and Xiaodan Zhu. 2020. [Speaker-aware BERT for multi-turn response selection in retrieval-based chatbots](#). In *CIKM '20: The 29th ACM International Conference on Information and Knowledge Management, Virtual Event, Ireland, October 19-23, 2020*, pages 2041–2044.
- Jia-Chen Gu, Zhen-Hua Ling, Xiaodan Zhu, and Quan Liu. 2019. [Dually interactive matching network for personalized 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 2019, Hong Kong, China, November 3-7, 2019*, pages 1845–1854. Association for Computational Linguistics.
- Jia-Chen Gu, Chongyang Tao, Zhen-Hua Ling, Can Xu, Xiubo Geng, and Daxin Jiang. 2021. [MPC-BERT: A pre-trained language model for multi-party conversation understanding](#). In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing, ACL/IJCNLP 2021, (Volume 1: Long Papers), Virtual Event, August 1-6, 2021*, pages 3682–3692. Association for Computational Linguistics.
- Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. 2016. [Deep residual learning for image recognition](#). In *2016 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2016, Las Vegas, NV, USA, June 27-30, 2016*, pages 770–778. IEEE Computer Society.
- Linmei Hu, Tianchi Yang, Chuan Shi, Houye Ji, and Xiaoli Li. 2019a. [Heterogeneous graph attention networks for semi-supervised short text classification](#). 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 2019, Hong Kong, China, November 3-7, 2019*, pages 4820–4829. Association for Computational Linguistics.
- Wenpeng Hu, Zhangming Chan, Bing Liu, Dongyan Zhao, Jinwen Ma, and Rui Yan. 2019b. [GSN: A graph-structured network for multi-party dialogues](#). In *Proceedings of the Twenty-Eighth International Joint Conference on Artificial Intelligence, IJCAI 2019, Macao, China, August 10-16, 2019*, pages 5010–5016.
- Ziniu Hu, Yuxiao Dong, Kuansan Wang, and Yizhou Sun. 2020. [Heterogeneous graph transformer](#). In *WWW '20: The Web Conference 2020, Taipei, Taiwan, April 20-24, 2020*, pages 2704–2710. ACM / IW3C2.
- Veton Kepuska and Gamal Bohouta. 2018. [Next-generation of virtual personal assistants \(microsoft cortana, apple siri, amazon alexa and google home\)](#). In *IEEE 8th Annual Computing and Communication Workshop and Conference, CCWC 2018, Las Vegas, NV, USA, January 8-10, 2018*, pages 99–103. IEEE.
- Thomas N. Kipf and Max Welling. 2017. [Semi-supervised classification with graph convolutional networks](#). In *5th International Conference on Learning Representations, ICLR 2017, Toulon, France, April 24-26, 2017, Conference Track Proceedings*. OpenReview.net.
- 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 2019, Hong Kong, China, November 3-7, 2019*, pages 1909–1919.
- Mike Lewis, Yinhan Liu, Naman Goyal, Marjan Ghazvininejad, Abdelrahman Mohamed, Omer Levy, Veselin Stoyanov, and Luke Zettlemoyer. 2020. [BART: denoising sequence-to-sequence pre-training for natural language generation, translation,](#)

- and comprehension. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics, ACL 2020, Online, July 5-10, 2020*, pages 7871–7880. Association for Computational Linguistics.
- Ilya Loshchilov and Frank Hutter. 2019. [Decoupled weight decay regularization](#). In *7th International Conference on Learning Representations, ICLR 2019, New Orleans, LA, USA, May 6-9, 2019*. OpenReview.net.
- 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 SIGDIAL 2015 Conference, The 16th Annual Meeting of the Special Interest Group on Discourse and Dialogue, 2-4 September 2015, Prague, Czech Republic*, pages 285–294.
- Hiroki Ouchi and Yuta Tsuboi. 2016. [Addressee and response selection for multi-party conversation](#). In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing, EMNLP 2016, Austin, Texas, USA, November 1-4, 2016*, pages 2133–2143.
- Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, and Ilya Sutskever. 2019. Language models are unsupervised multitask learners. *OpenAI blog*, 1(8):9.
- Michael Sejr Schlichtkrull, Thomas N. Kipf, Peter Bloem, Rianne van den Berg, Ivan Titov, and Max Welling. 2018. [Modeling relational data with graph convolutional networks](#). In *The Semantic Web - 15th International Conference, ESWC 2018, Heraklion, Crete, Greece, June 3-7, 2018, Proceedings*, volume 10843 of *Lecture Notes in Computer Science*, pages 593–607. Springer.
- Iulian Vlad Serban, Alessandro Sordani, Yoshua Bengio, Aaron C. Courville, and Joelle Pineau. 2016. [Building end-to-end dialogue systems using generative hierarchical neural network models](#). In *Proceedings of the Thirtieth AAAI Conference on Artificial Intelligence, February 12-17, 2016, Phoenix, Arizona, USA*, pages 3776–3784.
- Lifeng Shang, Zhengdong Lu, and Hang Li. 2015. [Neural responding machine for short-text conversation](#). In *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing of the Asian Federation of Natural Language Processing, ACL 2015, July 26-31, 2015, Beijing, China, Volume 1: Long Papers*, pages 1577–1586.
- Yizhou Sun and Jiawei Han. 2012. [Mining Heterogeneous Information Networks: Principles and Methodologies](#). Synthesis Lectures on Data Mining and Knowledge Discovery. Morgan & Claypool Publishers.
- Yizhou Sun, Jiawei Han, Xifeng Yan, Philip S. Yu, and Tianyi Wu. 2011. [Pathsim: Meta path-based top-k similarity search in heterogeneous information networks](#). *Proc. VLDB Endow.*, 4(11):992–1003.
- Yizhou Sun, Brandon Norrick, Jiawei Han, Xifeng Yan, Philip S. Yu, and Xiao Yu. 2012. [Integrating meta-path selection with user-guided object clustering in heterogeneous information networks](#). In *The 18th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, KDD '12, Beijing, China, August 12-16, 2012*, pages 1348–1356. ACM.
- Ilya Sutskever, Oriol Vinyals, and Quoc V. Le. 2014. [Sequence to sequence learning with neural networks](#). In *Advances in Neural Information Processing Systems 27: Annual Conference on Neural Information Processing Systems 2014, December 8-13 2014, Montreal, Quebec, Canada*, pages 3104–3112.
- Chongyang Tao, Wei Wu, Can Xu, Wenpeng Hu, Dongyan Zhao, and Rui Yan. 2019a. [Multi-representation fusion network for multi-turn response selection in retrieval-based chatbots](#). In *Proceedings of the Twelfth ACM International Conference on Web Search and Data Mining, WSDM 2019, Melbourne, VIC, Australia, February 11-15, 2019*, pages 267–275. ACM.
- Chongyang Tao, Wei Wu, Can Xu, Wenpeng Hu, Dongyan Zhao, and Rui Yan. 2019b. [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 Conference of the Association for Computational Linguistics, ACL 2019, Florence, Italy, July 28-August 2, 2019, Volume 1: Long Papers*, pages 1–11.
- Ming Tu, Guangtao Wang, Jing Huang, Yun Tang, Xiaodong He, and Bowen Zhou. 2019. [Multi-hop reading comprehension across multiple documents by reasoning over heterogeneous graphs](#). In *Proceedings of the 57th Conference of the Association for Computational Linguistics, ACL 2019, Florence, Italy, July 28- August 2, 2019, Volume 1: Long Papers*, pages 2704–2713. Association for Computational Linguistics.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin. 2017. [Attention is all you need](#). In *Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems 2017, December 4-9, 2017, Long Beach, CA, USA*, pages 5998–6008.
- Petar Velickovic, Guillem Cucurull, Arantxa Casanova, Adriana Romero, Pietro Liò, and Yoshua Bengio. 2018. [Graph attention networks](#). In *6th International Conference on Learning Representations, ICLR 2018, Vancouver, BC, Canada, April 30 - May 3, 2018, Conference Track Proceedings*. OpenReview.net.

- Danqing Wang, Pengfei Liu, Yining Zheng, Xipeng Qiu, and Xuanjing Huang. 2020a. [Heterogeneous graph neural networks for extractive document summarization](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics, ACL 2020, Online, July 5-10, 2020*, pages 6209–6219. Association for Computational Linguistics.
- Weishi Wang, Steven C. H. Hoi, and Shafiq R. Joty. 2020b. [Response selection for multi-party conversations with dynamic topic tracking](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing, EMNLP 2020, Online, November 16-20, 2020*, pages 6581–6591.
- Xiao Wang, Houye Ji, Chuan Shi, Bai Wang, Yanfang Ye, Peng Cui, and Philip S. Yu. 2019. [Heterogeneous graph attention network](#). In *The World Wide Web Conference, WWW 2019, San Francisco, CA, USA, May 13-17, 2019*, pages 2022–2032. ACM.
- Tsung-Hsien Wen, David Vandyke, Nikola Mrksic, Milica Gasic, Lina Maria Rojas-Barahona, Pei-Hao Su, Stefan Ultes, and Steve J. Young. 2017. [A network-based end-to-end trainable task-oriented dialogue system](#). In *Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics, EACL 2017, Valencia, Spain, April 3-7, 2017, Volume 1: Long Papers*, pages 438–449.
- Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Rémi Louf, Morgan Funtowicz, Joe Davison, Sam Shleifer, Patrick von Platen, Clara Ma, Yacine Jernite, Julien Plu, Canwen Xu, Teven Le Scao, Sylvain Gugger, Mariama Drame, Quentin Lhoest, and Alexander M. Rush. 2020. [Transformers: State-of-the-art natural language processing](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations, EMNLP 2020 - Demos, Online, November 16-20, 2020*, pages 38–45. Association for Computational Linguistics.
- Yu Wu, Wei Wu, Chen Xing, Ming Zhou, and Zhoujun Li. 2017. [Sequential matching network: A new architecture for multi-turn response selection in retrieval-based chatbots](#). In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics, ACL 2017, Vancouver, Canada, July 30 - August 4, Volume 1: Long Papers*, pages 496–505.
- Shaowei Yao, Tianming Wang, and Xiaojun Wan. 2020. [Heterogeneous graph transformer for graph-to-sequence learning](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics, ACL 2020, Online, July 5-10, 2020*, pages 7145–7154. Association for Computational Linguistics.
- Tom Young, Erik Cambria, Iti Chaturvedi, Hao Zhou, Subham Biswas, and Minlie Huang. 2018. [Augmenting end-to-end dialogue systems with commonsense knowledge](#). In *Proceedings of the Thirty-Second AAAI Conference on Artificial Intelligence, (AAAI-18), the 30th innovative Applications of Artificial Intelligence (IAAI-18), and the 8th AAAI Symposium on Educational Advances in Artificial Intelligence (EAAI-18), New Orleans, Louisiana, USA, February 2-7, 2018*, pages 4970–4977.
- Seongjun Yun, Minbyul Jeong, Raehyun Kim, Jaewoo Kang, and Hyunwoo J. Kim. 2019. [Graph transformer networks](#). In *Advances in Neural Information Processing Systems 32: Annual Conference on Neural Information Processing Systems 2019, NeurIPS 2019, December 8-14, 2019, Vancouver, BC, Canada*, pages 11960–11970.
- Chuxu Zhang, Dongjin Song, Chao Huang, Ananthram Swami, and Nitesh V. Chawla. 2019. [Heterogeneous graph neural network](#). In *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining, KDD 2019, Anchorage, AK, USA, August 4-8, 2019*, pages 793–803. ACM.
- Rui Zhang, Honglak Lee, Lazaros Polymenakos, and Dragomir R. Radev. 2018. [Addressee and response selection in multi-party conversations with speaker interaction rns](#). In *Proceedings of the Thirty-Second AAAI Conference on Artificial Intelligence, (AAAI-18), the 30th innovative Applications of Artificial Intelligence (IAAI-18), and the 8th AAAI Symposium on Educational Advances in Artificial Intelligence (EAAI-18), New Orleans, Louisiana, USA, February 2-7, 2018*, pages 5690–5697.
- Yizhe Zhang, Siqi 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, ACL 2020, Online, July 5-10, 2020*, pages 270–278. Association for Computational Linguistics.
- Li Zhou, Jianfeng Gao, Di Li, and Heung-Yeung Shum. 2020. [The design and implementation of xiaoice, an empathetic social chatbot](#). *Comput. Linguistics*, 46(1):53–93.
- 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, ACL 2018, Melbourne, Australia, July 15-20, 2018, Volume 1: Long Papers*, pages 1118–1127.