

Structure-Aware Abstractive Conversation Summarization via Discourse and Action Graphs

Jiaao Chen

School of Interactive Computing
Georgia Institute of Technology
jiaaochen@gatech.edu

Diyi Yang

School of Interactive Computing
Georgia Institute of Technology
dyang888@gatech.edu

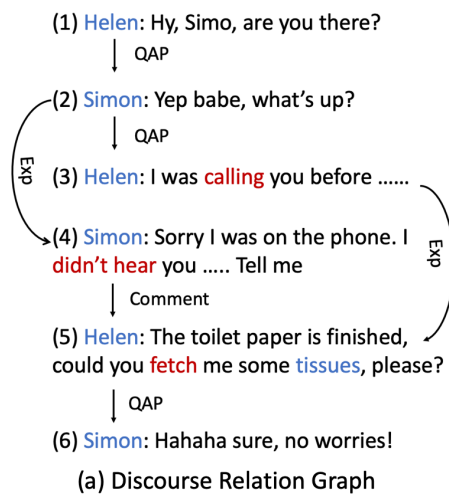
Abstract

Abstractive conversation summarization has received much attention recently. However, these generated summaries often suffer from insufficient, redundant, or incorrect content, largely due to the unstructured and complex characteristics of human-human interactions. To this end, we propose to explicitly model the rich structures in conversations for more precise and accurate conversation summarization, by first incorporating discourse relations between utterances and action triples (“WHO-DOING-WHAT”) in utterances through structured graphs to better encode conversations, and then designing a multi-granularity decoder to generate summaries by combining all levels of information. Experiments show that our proposed models outperform state-of-the-art methods and generalize well in other domains in terms of both automatic evaluations and human judgments. We have publicly released our code at <https://github.com/GT-SALT/Structure-Aware-BART>.

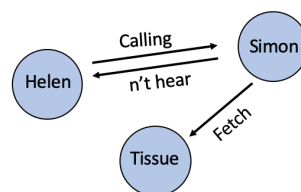
1 Introduction

Online interaction has become an indispensable component of everyday life and people are increasingly using textual conversations to exchange ideas, make plans, and share information. However, it is time-consuming to recap and grasp all the core content within every complex conversation (Gao et al., 2020; Feng et al., 2020). As a result, how to organize massive everyday interactions into natural, concise, and informative text, i.e., abstractive conversation summarization, starts to gain importance.

Significant progress has been made on abstractive summarization for structured document via pointer generator (See et al., 2017), reinforcement methods (Paulus et al., 2018; Huang et al., 2020a) and pre-trained models (Liu and Lapata, 2019; Lewis et al., 2020; Zhang et al., 2019). Despite the huge success, it is challenging to directly apply document models to summarize conversations, due to



(a) Discourse Relation Graph



(b) Action Graph

Figure 1: An example of discourse relation graph (a) and action graph (b) from one conversation in SAM-Sum (Gliwa et al., 2019). The annotated summary is *Simon was on the phone before, so he didn't here Helen calling. Simon will fetch Helen some tissues.*

a set of inherent differences between conversations and documents (Gliwa et al., 2019). First, speaker interruptions like repetitions, false-starts, and hesitations are frequent in conversations (Sacks et al., 1978), and key information resides in different portions of a conversation. These unstructured properties pose challenges for models to focus on salient contents that are necessary for generating both abstractive and informative summaries. Second, there is more than one speaker in conversations and people interact with each other in different language styles (Zhu et al., 2020b). The complex interactions among multiple speakers make it harder for mod-

els to identify and associate speakers with correct actions so as to generate factual summaries.

In order to summarize the unstructured and complex conversations, a growing body of research has been conducted, such as transferring document summarization methods to conversation settings (Shang et al., 2018; Gliwa et al., 2019), adopting hierarchical models (Zhao et al., 2019; Zhu et al., 2020b), or incorporating conversation structures like topic segmentation (Liu et al., 2019b; Li et al., 2019; Chen and Yang, 2020), dialogue acts (Goo and Chen, 2018), and conversation stages (Chen and Yang, 2020). However, current approaches still face challenges in terms of succinctness and faithfulness, as most prior studies (i) fail to explicitly model dependencies between utterances which can help identify salient portions of conversations (Bui et al., 2009), and (ii) lack structured representations (Huang et al., 2020a) to learn the associations between speakers, actions and events. We argue that these rich linguistic structures associated with conversations are key components towards generating abstractive and factual conversation summaries.

To this end, we present a structure-aware sequence-to-sequence model, in which we equip abstractive conversation summarization models with rich conversation structures through two types of graphs: **discourse relation graph** and **action graph**. Discourse relation graphs are constructed based on dependency-based discourse relations (Kirschner et al., 2012; Stone et al., 2013; Asher et al., 2016; Qin et al., 2017) between intertwined utterances, where each Elementary Discourse Unit (EDU) is one single utterance and they are linked through 16 different types of relations (Asher et al., 2016). As shown in Figure 1(a), highly related utterances are linked based on discourse relations like *Question Answer Pairs*, *Comment* and *Explanation*. Explicitly modeling these utterances relations in conversations can aid models in recognizing key content for succinct and informative summarization. Action graphs are constructed as the “WHO-DOING-WHAT” triplets in conversations which express socially situated identities and activities (Gee, 2014). For instance, in Figure 1(b), the action graph provides explicit information between *Simon*, *fetch*, and *tissues* for the utterance *it is Simon who will fetch the tissues*, making models less likely to generate summaries with wrong references (e.g., *Helen will fetch the tissues*).

To sum up, our contributions are: (1) We pro-

pose to utilize discourse relation graphs and action graphs to better encode conversations for conversation summarization. (2) We design structure-aware sequence-to-sequence models to combine these structured graphs and generate summaries with the help of a novel multi-granularity decoder. (3) We demonstrate the effectiveness of our proposed methods through experiments on a large-scale conversation summarization dataset, SAM-Sum (Gliwa et al., 2019). (4) We further show that our structure-aware models can generalize well in new domains such as debate summarization.

2 Related Work

Document Summarization Compared to extractive document summarization (Gupta and Lehal, 2010; Narayan et al., 2018; Liu and Lapata, 2019), abstractive document summarization is generally considered more challenging and has received more attention. Various methods have been designed to tackle abstractive document summarization like sequence-to-sequence models (Rush et al., 2015), pointer generators (See et al., 2017), reinforcement learning methods (Paulus et al., 2018; Huang et al., 2020a) and pre-trained models (Lewis et al., 2020; Zhang et al., 2019). To generate faithful abstractive document summaries (Maynez et al., 2020), graph-based models were introduced recently such as extracting entity types (Fernandes et al., 2018; Fan et al., 2019), leveraging knowledge graphs (Huang et al., 2020a; Zhu et al., 2020a) or designing extra fact correction modules (Dong et al., 2020). Inspired by these graph-based methods, we also construct action graphs for generating more factual conversation summaries.

Conversation Summarization Extractive dialogue summarization (Murray et al., 2005) has been studied extensively via statistical machine learning methods such as skip-chain CRFs (Galley, 2006), SVM with LDA models (Wang and Cardie, 2013), and multi-sentence compression algorithms (Shang et al., 2018). Such methods struggled with generating succinct, fluent, and natural summaries, especially when the key information needs to be aggregated from multiple first-person point-of-view utterances (Song et al., 2020). Abstractive conversation summarization overcomes these issues by designing hierarchical models (Zhao et al., 2019; Zhu et al., 2020b), incorporating commonsense knowledge (Feng et al., 2020), or leveraging conversational structures like dialogue acts (Goo and Chen,

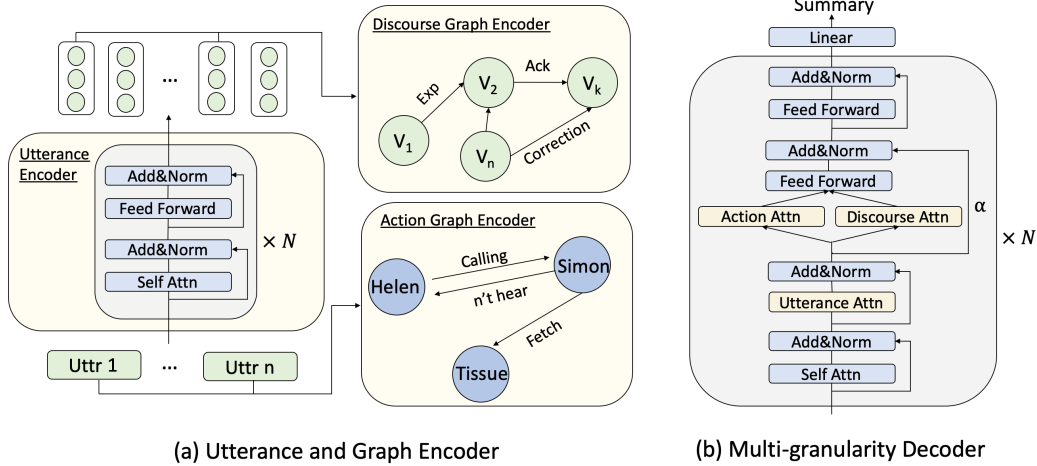


Figure 2: Model architecture. Each utterance is encoded via transformer encoder; discourse relation graphs and action graphs are encoded through Graph Attention Networks (a). The multi-granularity decoder (b) then generates summaries based on all levels of encoded information including utterances, action graphs, and discourse graphs.

2018), key point sequences (Liu et al., 2019a), topic segments (Liu et al., 2019b; Li et al., 2019) and stage developments (Chen and Yang, 2020). Some recent research has also utilized discourse relations as input features in classifiers to detect important content in conversations (Murray et al., 2006; Bui et al., 2009; Qin et al., 2017). However, current models still have not explicitly utilized the dependencies between different utterances, making models hard to leverage long-range dependencies and utilize these salient utterances. Moreover, less attention has been paid to identify the actions of different speakers and how they interact with or refer to each other, leading to unfaithful summarization with incorrect references or wrong reasoning (Gliwa et al., 2019). To fill these gaps, we propose to explicitly model actions within utterances, and relations between utterances in conversations in a structured way, by using discourse relation graphs and action graphs and further combining these through relational graph encoders and multi-granularity decoders for abstractive conversation summarization.

3 Methods

To generate abstractive and factual summaries from unstructured conversations, we propose to model structural signals in conversations by first constructing discourse relation graphs and action graphs (Section 3.1), and then encoding the graphs together with conversations (Section 3.2) as well as incorporating these different levels of information in the decoding stage through a multi-granularity

decoder (Section 3.3) to summarize given conversations. The overall architecture is shown in Figure 2.

3.1 Structured Graph Construction

This section describes how to construct the discourse relation graphs and action graphs. Formally, for a given conversation $C = \{u_0, \dots, u_m\}$ with m utterances, we construct *discourse relation graph* $\mathcal{G}^D = (\mathbf{V}^D, \mathbf{E}^D)$, where \mathbf{V}^D is the set of nodes representing Elementary Discourse Units (EDUs), and \mathbf{E}^D is the adjacent matrix that describes the relations between EDUs, and *action graph* $\mathcal{G}^A = (\mathbf{V}^A, \mathbf{E}^A)$, where \mathbf{V}^A is the set of nodes representing “WHO”, “DOING” and “WHAT” arguments, and \mathbf{E}^A is the adjacent matrix to link “WHO-DOING-WHAT” triples.

Discourse Relation Graph Utterances from different speakers do not occur in isolation; instead, they are related within the context of discourse (Murray et al., 2006; Qin et al., 2017), which has been shown effective for dialogue understanding like identifying the decisions in multi-party dialogues (Bui et al., 2009) and detecting salient content in email conversations (McKeown et al., 2007). Although current attention-based neural models are supposed to, or might implicitly, learn certain relations between utterances, they often struggle to focus on many informative utterances (Chen and Yang, 2020; Song et al., 2020) and fail to address long-range dependencies (Xu et al., 2020), especially when there are frequent interruptions. As a result, explicitly incorporating the discourse relations will help neural summarization models better

encode the unstructured conversations and concentrate on the most salient utterances to generate more informative and less redundant summaries.

To do so, we view each utterance as an EDU and use the discourse relation types defined in Asher et al. (2016). We first pre-train a discourse parsing model (Shi and Huang, 2019) on a human-annotated multiparty dialogue corpus (Asher et al., 2016), with 0.775 F1 score on link predictions and 0.557 F1 score on relation classifications, which are comparable to the state-of-the-art results (Shi and Huang, 2019). We then utilize this pre-trained parser to predict the discourse relations within conversations in our SAMSum corpus (Gliwa et al., 2019).

After predictions, there are 138,554 edges identified in total and 8.48 edges per conversation. The distribution of these predicted discourse relation types is: Comment (19.3%), Clarification Question (15.2%), Elaboration (2.3%), Acknowledgement (8.4%), Continuation (10.1%), Explanation (2.8%), Conditional (0.2%), Question Answer Pair (21.5%), Alternation (0.3%), Q-Elab (2.5%), Result (5.5%), Background (0.4%), Narration (0.4%), Correction (0.4%), Parallel (0.9%), and Contrast (1.0%). Then for each conversation, we construct a discourse relation graph $\mathcal{G}^D = (\mathbf{V}^D, \mathbf{E}^D)$, where $\mathbf{V}^D[k]$ represents the k -th utterance. $\mathbf{E}^D[i][j] = r$ if there is a link from the i -th utterance to the j -th one with discourse relation r .

Action Graph The “*who-doing-what*” triples from utterances can provide explicit visualizations of speakers and their actions, the key to understanding concrete details happened in conversations (Moser, 2001; Gee, 2014; Sacks et al., 1978). Simply relying on neural models to identify this information from conversations often fail to produce factual characterizations of concrete details happened (Cao et al., 2018; Huang et al., 2020a). To this end, we extract “WHO-DOING-WHAT” triples from utterances and construct action graphs for conversation summarization (Chen et al., 2019; Huang et al., 2020b,a). Specifically, we first transform the first-person point-of-view utterances to its third-person point-of-view forms based on simple rules: (i) substituting first/second-person pronouns with the names of current speaker or surrounding speakers and (ii) replacing third-person pronouns based on coreference clusters in conversations detected by the Stanford CoreNLP (Manning et al., 2014). For example, an utterance “*I’ll bring it to you to-*

morrow” from Amanda to Jerry will be transformed into “*Amanda’ll bring cakes to Jerry tomorrow*”. Then we extract “WHO-DOING-WHAT” (subject-predicate-object) triples from transformed conversations using the open information extraction (OpenIE) systems¹ (Angeli et al., 2015). We then construct the Action Graph $\mathcal{G}^A = (\mathbf{V}^A, \mathbf{E}^A)$ from the extracted triples by taking arguments (“WHO”, “DOING”, or “WHAT”) as nodes in \mathbf{V}^A , and connect them with edge $\mathbf{E}^A[i][j] = 1$ if they are adjacent in one “WHO-DOING-WHAT” triple.

3.2 Encoder

Given a conversation and its corresponding discourse relation graph and action graph, we utilize an utterance encoder and two graph encoders, to obtain its hidden representations shown in Figure 2(a).

3.2.1 Utterance Encoder

We initialize our utterance encoder $F_U(\cdot)$ with a pre-trained encoder, i.e., BART-base (Lewis et al., 2020), and encode tokens $\{x_{i,0}, \dots, x_{i,l}\}$ in an utterance \mathbf{u}_i into its hidden representation:

$$\{h_{i,0}^U, \dots, h_{i,l}^U\} = F_U(\{x_{i,0}, \dots, x_{i,l}\}) \quad (1)$$

Here we add a special token $x_{i,0} = \langle S \rangle$ at the beginning of each utterance to represent it.

3.2.2 Graph Encoder

Node Initialization For *discourse relation graph*, we employ the output embeddings of the special tokens $x_{i,0}$ from the utterance encoder, i.e., $h_{i,0}^U$, to initialize the i -th node v_i^D in \mathcal{G}^D . We use a one-hot embedding layer to encode the relations $\mathbf{E}^D[i][j] = e_{i,j}^D$ between utterance i and j . For *action graph*, we first utilize $F_U(\cdot)$ to encode each token in nodes v_i^A and then average their output embeddings as their initial representations.

Structured Graph Attention Network Based on Graph Attention Network (Veličković et al., 2018), we utilize these relations between nodes to encode each node v_i^D in \mathcal{G}^D or v_i^A in \mathcal{G}^A through:

$$\alpha_{ij} = \frac{\exp(\sigma(\mathbf{a}^T [\mathbf{W}v_i \parallel \mathbf{W}v_j \parallel \mathbf{W}_e e_{i,j}]))}{\sum_{k \in \mathcal{N}_i} \exp(\sigma(\mathbf{a}^T [\mathbf{W}v_i \parallel \mathbf{W}v_k \parallel \mathbf{W}_e e_{i,k}]))}$$

$$h_i = \sigma\left(\sum_{j \in \mathcal{N}_i} \alpha_{ij} \mathbf{W}v_j\right)$$

¹<https://github.com/philipperemy/Stanford-OpenIE-Python>

Dataset	Split	# Conv	# Participants	# Turns	# Discourse Edges	# Action Triples
SAMSum	Train	14732	2.40	11.17	8.47	6.72
	Val	818	2.39	10.83	8.34	6.48
	Test	819	2.36	11.25	8.63	6.81
ADSC	Full	45	2.00	7.51	6.51	37.20

Table 1: Statistics of the used datasets, including the total number of conversations (# Conv), the average number of participants, turns, discourse edges and action triples per conversation.

\mathbf{W} , \mathbf{W}_e and \mathbf{a} are trainable parameters. $[\cdot\|\cdot]$ denotes the concatenation of two vectors. σ is the activation function, \mathcal{N}_i is the set containing node- i 's neighbours in \mathcal{G} .

Through two graph encoders $F_D(\cdot, \cdot)$ and $F_A(\cdot, \cdot)$, we then obtain the hidden representations of these nodes as:

$$\{h_0^D, \dots, h_m^D\} = F_D(\{v_0^D, \dots, v_m^D\}, \mathbf{E}^D) \quad (2)$$

$$\{h_0^A, \dots, h_n^A\} = F_A(\{x_0^A, \dots, x_n^A\}, \mathbf{E}^A) \quad (3)$$

3.3 Multi-Granularity Decoder

Different levels of encoded representations are then aggregated via our multi-granularity decoder to generate summaries as shown in Figure 2(b). With $s - 1$ previously generated tokens y_1, \dots, y_{s-1} , our decoder $G(\cdot)$ predicts the l -th token via:

$$\hat{y} = G(y_{1:s-1}, F_U(\mathbf{C}), F_D(\mathcal{G}^D), F_A(\mathcal{G}^A)) \quad (4)$$

$$P(\tilde{y}_s | y_{<s}, \mathbf{C}, \mathcal{G}^D, \mathcal{G}^A) = \text{Softmax}(W_p \hat{y}) \quad (5)$$

To better incorporate the information in constructed graphs, different from the traditional pre-trained BART model (Lewis et al., 2020), we improve the BART transformer decoder with two extra cross attentions (Discourse Attention and Action Attention) added to each decoder layer, which attends to the encoded node representations in discourse relation graphs and action graphs.

In each decoder layer, after performing the original cross attentions over every token in utterances $\{h_{i,0:l}^U\}$ and getting the utterance-attended representation x^U , multi-granularity decoder then conducts cross attentions over nodes $\{h_{0:m}^D\}$ and $\{h_{0:n}^A\}$ that are encoded from graph encoders in parallel, to obtain the discourse-attended representation x^D and action-attended representation x^A . These two attended vectors are then combined into a structure-aware representation x^S , through a feed-forward network for further forward passing in the decoder.

To alleviate the negative impact of randomly initialized graph encoders and cross attentions over graphs on pre-trained BART decoders at

early stages and accelerate the learning of newly-introduced modules during training, we apply ReZero (Bachlechner et al., 2020) to the residual connection after attending to graphs in each decoder layer:

$$\tilde{x}^S = x^U + \alpha x^S \quad (6)$$

where α is one trainable parameter instead of a fixed value 1, which modulates updates from cross attentions over graphs.

Training During training, we seek to minimize the cross entropy and use the teacher-forcing strategy (Bengio et al., 2015):

$$\mathcal{L} = - \sum \log P(\tilde{y}_l | y_{<l}, \mathbf{C}, \mathcal{G}^D, \mathcal{G}^A) \quad (7)$$

4 Experiments

4.1 Datasets

We trained and evaluated our models on a conversation summarization dataset SAMSum (Gliwa et al., 2019) covering messenger-like conversations about daily topics, such as arranging meetings and discussing events. We also showed the generalizability of our models on the Argumentative Dialogue Summary Corpus (ADSC) (Misra et al., 2015), a debate summarization corpus. The data statistics of two datasets were shown in Table 1, with the discourse relation types distributions in the Appendix.

4.2 Baselines

We compare our methods with several baselines:

- **Pointer Generator** (See et al., 2017): We followed the settings in Gliwa et al. (2019) and used special tokens to separate each utterance.
- **Transformer** (Vaswani et al., 2017): We trained transformer seq2seq models following the OpenNMT (Klein et al., 2017).
- **D-HGN** (Feng et al., 2020) incorporated commonsense knowledge from ConceptNet (Liu and Singh, 2004) for dialogue summarization.

Model	ROUGE-1			ROUGE-2			ROUGE-L		
	F	P	R	F	P	R	F	P	R
Pointer Generator (See et al., 2017)	40.08	-	-	15.28	-	-	36.63	-	-
Transformer (Vaswani et al., 2017)	37.27	-	-	10.76	-	-	32.73	-	-
D-HGN (Feng et al., 2020)	42.03	-	-	18.07	-	-	39.56	-	-
Multi-view Seq2Seq (Chen and Yang, 2020)	45.56	52.13	44.68	22.30	25.58	22.03	44.70	50.82	43.29
BART (Lewis et al., 2020)	45.15	49.58	45.97	21.66	23.95	22.16	44.46	48.92	44.26
S-BART w. Discourse †	45.89	51.34	45.87	22.50	25.26	22.33	44.83	49.93	44.17
S-BART w. Action †	45.67	50.25	46.44	22.39	24.70	22.96	44.86	49.29	44.75
S-BART w. Discourse&Action †	46.07	51.13	46.24	22.60	25.11	22.81	45.00	49.82	44.47

Table 2: ROUGE-1, ROUGE-2 and ROUGE-L scores for different models on the SAMSum Corpus test set. Results are averaged over three random runs. † means our methods. We performed Pitman’s permutation test (Dror et al., 2018) and found that *S-BART w. Discourse&Action* significantly outperformed the base *BART* ($p < 0.05$).

Model	ROUGE-1			ROUGE-2			ROUGE-L		
	F	P	R	F	P	R	F	P	R
BART (Lewis et al., 2020)	20.90	51.71	13.53	5.04	12.46	3.24	21.23	56.29	13.54
S-BART w. Discourse †	22.42	54.13	14.60	5.58	13.83	3.61	22.16	51.88	14.45
S-BART w. Action †	30.91	85.42	19.12	20.64	56.31	12.78	35.30	85.51	22.58
S-BART w. Discourse&Action †	34.74	84.99	22.20	23.86	58.08	15.24	38.69	83.81	25.51

Table 3: ROUGE-1, ROUGE-2 and ROUGE-L scores on the out-of-domain ADSC corpus using different models trained on SAMSum Corpus. † means our methods.

- **BART** (Lewis et al., 2020): We utilized BART², and separated utterances by a special token.
- **Multi-View Seq2Seq** (Chen and Yang, 2020) utilized topic and stage views on top of BART for summarizing conversations. Here we implemented it based on BART-base models.

4.3 Implementation Details

We used the BART-base model to initialize our sequence-to-sequence model for training in all experiments. For parameters in the original BART encoder/decoder, we followed the default settings and set the learning rate $3e-5$ with 120 warm-up steps. For graph encoders, we set the number of hidden dimensions as 768, the number of attention heads as 2, the number of layers as 2, and the dropout rate as 0.2. For graph cross attentions added to BART decoder layers, we set the number of attention heads as 2. The weights α in ReZero residual connections were initialized with 1. The learning rate for parameters in newly added modules was $3e-4$ with 60 warm-up steps. All experiments were performed on GeForce RTX 2080Ti (11GB memory).

4.4 Results on In-Domain Corpus

Automatic Evaluation We evaluated all the models with the widely used automatic metric,

²The version on 10/7 in https://huggingface.co/transformers/model_doc/bart.html

Models	Fac.	Suc.	Inf.
Ground Truth	4.29	4.40	4.06
BART	3.90	4.13	3.74
S-BART w. Discourse	4.11	4.42	3.98
S-BART w. Action	4.17	4.29	3.95
S-BART w. Discourse&Action	4.19	4.41	3.91

Table 4: Human evaluation on **F**actualness, **S**uccinctness, **I**nformativeness. All model variants of *S-BART* received significantly higher ratings than *BART* (student t-test, $p < 0.05$).

ROUGE scores (Lin and Och, 2004)³, and reported ROUGE-1, ROUGE-2, and ROUGE-L in Table 2. We found that, compared to simple sequence-to-sequence models (*Pointer Generator* and *Transformer*), incorporating extra information such as commonsense knowledge from ConceptNet (*D-HGN*) increased the ROUGE metrics. When equipped with pre-trained models and simple conversation structures such as topics and conversation stages, *Multi-View Seq2Seq* boosted ROUGE scores. Incorporating discourse relation graphs or action graphs helped the performances of summarization, suggesting the effectiveness of explicitly modeling relations between utterances and the associations between speakers and actions within utterances. Combining two different structured graphs produced better ROUGE scores compared to previous state-of-the-art methods and our base models,

³We followed fairseq and used <https://github.com/pltrdy/rouge> to calculate ROUGE scores. Note that different tools may result in different ROUGE scores.

with an increase of 2.0% on ROUGE-1, 4.3% on ROUGE-2, and 1.2% on ROUGE-L compared to our base model, *BART*. This indicates that, our structure-aware models with discourse and action graphs could help abstractive conversation summarization, and these two graphs complemented each other in generating better summaries.

Human Evaluation We conducted human evaluation to qualitatively evaluate the generated summaries. Specifically, we asked annotators from Amazon Mechanical Turk to score a set of randomly sampled 100 generated summaries from ground-truth, *BART* and our structured models, using a Likert scale from 1 (worst) to 5 (best) in terms of **factualness** (e.g., associates actions with the right actors), **succinctness** (e.g., does not contain redundant information), and **informativeness** (e.g., covers the most important content) (Feng et al., 2020; Huang et al., 2020a). To increase annotation quality, we required turkers to have a 98% approval rate and at least 10,000 approved tasks for their previous work. Each message was rated by three workers. The scores for each summary were averaged. The Intra-Class Correlation was 0.543, showing moderate agreement (Koo and Li, 2016).

As shown in Table 4, *S-BART* that utilized structured information from discourse relation graphs and action graphs generated significantly better summaries with respect to factualness, succinctness, and informativeness. This might be because that the incorporation of structured information such as discourse relations helped *S-BART* to recognize the salient parts in conversations, and thus improve the succinctness and informativeness over *BART*. Modeling the connections between speakers and actions greatly helped generate more factual summaries than the baselines, e.g., with an increase of 0.27 from *BART* to *S-BART w. Action*.

4.5 Results on Out-Of-Domain Corpus

To investigate the generalizability of our structure-aware models, we then tested the *S-BART* model trained on SAMSum corpus directly on the debate summarization domain (ADSC Corpus (Misra et al., 2015)) in a zero-shot setting. Besides the differences in topics, utterances in debate conversations were generally longer and include more action triples (37.20 vs 6.81 as shown in Table 1) and fewer participants. The distribution of discourse relation types also differed a lot across different

Graph Types	R-1	R-2	R-L
S-BART w. Discourse Graph	45.89	22.50	44.83
S-BART w. Random Graph	45.28	21.80	44.30

Table 5: ROUGE-1, ROUGE-2 and ROUGE-L scores of S-BART with either the constructed discourse relation graphs or random graphs. Results are averaged over three random runs.

Combination Strategy	R-1	R-2	R-L
Parallel	46.07	22.60	45.00
Sequential (discourse, action)	45.40	22.14	44.67
Sequential (action, discourse)	45.62	22.41	44.62

Table 6: ROUGE-1, ROUGE-2 and ROUGE-L scores of S-BART models using different ways to combine discourse relation graphs and action graphs. Results are averaged over three random runs.

domains⁴ (e.g., more *Contrast* in debates (19.5%) than in daily conversations (1.0%)).

As shown in Table 3, our single graph models *S-BART w. Discourse* and *S-BART w. Action* boosted ROUGE scores compared to *BART*, suggesting that utilizing structures can also increase the generalizability of conversation summarization methods. However, contrary to in-domain results in Table 2, action graphs led to much more gains than discourse graphs. This indicated that when domain shifts, action triples were most robust in terms of zero-shot setups; differences in discourse relation distributions could limit such generalization. Consistent with in-domain scenarios, our *S-BART w. Discourse&Action* achieved better results, with an increase of 66.2% on ROUGE-1, 373.4% on ROUGE-2, and 82.2% on ROUGE-L over *BART*.

4.6 Ablation Studies

This part conducted ablation studies to show the effectiveness of structured graphs in our *S-BART*.

The Quality of Discourse Relation Graphs We showed how the quality of discourse relation graphs affected the performances of conversation summarization in Table 5. Specifically, we compared the ROUGE scores of S-BART using our constructed discourse relation graphs (*S-BART w. Discourse Graph*) and S-BART using randomly generated discourse relation graphs *S-BART w. Random Graph* where both connections between nodes and relation types were randomized. The number of edges in two graphs was kept the same. We found that S-BART with our discourse graphs outperformed

⁴The detailed distributions were shown in the Appendix.

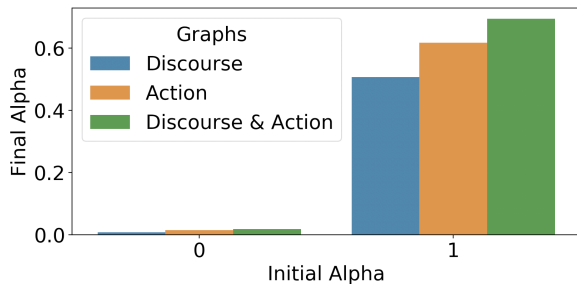


Figure 3: Averaged α over decoder layers in the trained S-BART models using different graphs

models with random graphs, indicating the effectiveness of the constructed discourse relation graphs and the importance of their qualities.

Different Ways to Combine Graphs We experimented with different ways to combine discourse relation graphs and action graphs in our *S-BART w. Discourse & Action*, and presented the results in Table 6. Here, **parallel** strategy performed cross attentions on different graphs separately and then combined the attended results with feed-forward networks as discussed in Section 3.3; **sequential** strategy performed cross attentions on two graphs in a specific order (from discourse relation graphs to actions graphs, or vice versa). We found that the parallel strategy showed better performances and the sequential ones did not introduce gains compared to S-BART with single graphs. This demonstrates that discourse relation graphs and action graphs were both important and provided different signals for abstractive conversation summarization.

Visualizing ReZero Weights We further tested our structure-aware BART with two ReZero settings: (i) initializing α from 0, (ii) initializing α from 1, and found initializing α from 1 would bring in more performance gains (see Appendix). We then visualized the average α over different decoder layers after training in Figure 3, and observed that (i) when α was initialized with 1, the final α was much larger than the setting where α was initialized with 0, which might be because randomly initialized modules barely received supervisions at early stages and therefore contributes less to BART. (ii) Compared to discourse graphs, action graphs received higher α weights after training in both initializing settings, suggesting that the information from structured action graphs might be harder for the end-to-end BART models to capture. (iii) Utilizing both graphs spontaneously led to higher

Conversations	# Num	# Dis.	# Act.
Test Set	819	8.63	6.81
Similar	373	8.31	6.36
Increase	208	9.13	7.40
Challenging	160	9.58	7.85

Table 7: The total number of examples, average number of **Discourse** edges and **Action** triples in different set of conversations in the SAMSUM test set.

ReZero weights, further validating the effectiveness of combining discourse relation graphs and action graphs and their complementary properties.

4.7 Error Analyses

To inspect when our summarization models could help the conversations summarization, we visualized the average number of discourse edges and the average number of action triples in three sets of conversations in Table 7: (i) **Similar**: examples where *S-BART* generated similar ROUGE scores (the differences were less than 0.1) compared to *BART*; (ii) **Increase**: examples where *S-BART* resulted in higher ROUGE scores (the differences were larger than 1.0) compared to *BART*; (iii) **Challenging**: examples where both *S-BART* and *BART* showed low ROUGE scores (ROUGE-1 < 20.0, ROUGE-2 < 10.0, ROUGE-L < 10.0).

When the structures in conversations were simpler (fewer discourse edges and fewer action triples than the average), *BART* showed similar performance as *S-BART*. As the structures of conversations become more complex with more discourse relations and more action mentions, *S-BART* outperformed *BART* as it explicitly incorporated these structured graphs. However, both *BART* and *S-BART* struggled when there were much more interactions beyond certain thresholds, calling for better mechanisms to model structures in conversations for generating better summaries.

5 Conclusion

In this work, we introduced a structure-aware sequence-to-sequence model for abstractive conversation summarization by incorporating discourse relations between utterances, and the connections between speakers and actions within utterances. Experiments and ablation studies on SAMSUM corpus showed the effectiveness of these structured graphs in aiding the task of conversation summarization via both quantitative and qualitative eval-

uation metrics. Results in zero-shot settings on ADCS Corpus further demonstrated the generalizability of our structure-aware models. In the future, we plan to extend our current conversation summarization models for various application domains such as emails, debates, and podcasts, and in conversations that might involve longer utterances and more participants in an unsynchronized way.

Acknowledgement

We would like to thank the anonymous reviewers for their helpful comments, and the members of Georgia Tech SALT group for their feedback. This work is supported in part by grants from Google, Amazon and Salesforce.

References

- Gabor Angeli, Melvin Jose Johnson Premkumar, and Christopher D. Manning. 2015. [Leveraging linguistic structure for open domain information extraction](#). In *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 344–354, Beijing, China. Association for Computational Linguistics.
- Nicholas Asher, Julie Hunter, Mathieu Morey, Benamara Farah, and Stergos Afantenos. 2016. Discourse structure and dialogue acts in multiparty dialogue: the stac corpus. In *Proceedings of the Tenth International Conference on Language Resources and Evaluation (LREC’16)*, pages 2721–2727.
- Thomas Bachlechner, Bodhisattwa Prasad Majumder, Huanru Henry Mao, Garrison W Cottrell, and Julian McAuley. 2020. Rezero is all you need: Fast convergence at large depth. *arXiv preprint arXiv:2003.04887*.
- Samy Bengio, Oriol Vinyals, Navdeep Jaitly, and Noam Shazeer. 2015. Scheduled sampling for sequence prediction with recurrent neural networks. In *Advances in Neural Information Processing Systems*, pages 1171–1179.
- Trung Bui, Matthew Frampton, John Dowding, and Stanley Peters. 2009. [Extracting decisions from multi-party dialogue using directed graphical models and semantic similarity](#). In *Proceedings of the SIGDIAL 2009 Conference*, pages 235–243, London, UK. Association for Computational Linguistics.
- Ziqiang Cao, Furu Wei, W. Li, and Sujian Li. 2018. Faithful to the original: Fact aware neural abstractive summarization. In *AAAI*.
- Jiaao Chen, Jianshu Chen, and Zhou Yu. 2019. Incorporating structured commonsense knowledge in story completion. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, pages 6244–6251.
- Jiaao Chen and Diyi Yang. 2020. [Multi-view sequence-to-sequence models with conversational structure for abstractive dialogue summarization](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 4106–4118, Online. Association for Computational Linguistics.
- Yue Dong, Shuohang Wang, Zhe Gan, Yu Cheng, Jackie Chi Kit Cheung, and Jingjing Liu. 2020. [Multi-fact correction in abstractive text summarization](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 9320–9331, Online. Association for Computational Linguistics.
- Rotem Dror, Gili Baumer, Segev Shlomov, and Roi Reichart. 2018. The hitchhiker’s guide to testing statistical significance in natural language processing. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1383–1392.
- Angela Fan, Claire Gardent, Chloé Braud, and Antoine Bordes. 2019. Using local knowledge graph construction to scale seq2seq models to multi-document inputs. 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 4177–4187.
- Xiachong Feng, Xiaocheng Feng, Bing Qin, and Ting Liu. 2020. Incorporating commonsense knowledge into abstractive dialogue summarization via heterogeneous graph networks. *arXiv preprint arXiv:2010.10044*.
- Patrick Fernandes, Miltiadis Allamanis, and Marc Brockschmidt. 2018. Structured neural summarization. In *International Conference on Learning Representations*.
- Michel Galley. 2006. A skip-chain conditional random field for ranking meeting utterances by importance. In *Proceedings of the 2006 Conference on Empirical Methods in Natural Language Processing*, pages 364–372.
- Shen Gao, Xiuying Chen, Zhaochun Ren, Dongyan Zhao, and Rui Yan. 2020. From standard summarization to new tasks and beyond: Summarization with manifold information. *arXiv preprint arXiv:2005.04684*.
- James Paul Gee. 2014. *An introduction to discourse analysis: Theory and method*. Routledge.
- Bogdan Gliwa, Iwona Mochol, Maciej Biesek, and Aleksander Wawer. 2019. [SAMSUM corpus: A human-annotated dialogue dataset for abstractive summarization](#). In *Proceedings of the 2nd Workshop*

- on *New Frontiers in Summarization*, pages 70–79, Hong Kong, China. Association for Computational Linguistics.
- Chih-Wen Goo and Yun-Nung Chen. 2018. [Abstractive dialogue summarization with sentence-gated modeling optimized by dialogue acts](#). *2018 IEEE Spoken Language Technology Workshop (SLT)*.
- Vishal Gupta and Gurpreet Singh Lehal. 2010. A survey of text summarization extractive techniques. *Journal of emerging technologies in web intelligence*, 2(3):258–268.
- Luyang Huang, Lingfei Wu, and Lu Wang. 2020a. [Knowledge graph-augmented abstractive summarization with semantic-driven cloze reward](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 5094–5107, Online. Association for Computational Linguistics.
- Shanshan Huang, Kenny Q. Zhu, Qianzi Liao, Libin Shen, and Yingong Zhao. 2020b. [Enhanced story representation by conceptnet for predicting story endings](#). In *Proceedings of the 29th ACM International Conference on Information Knowledge Management, CIKM '20*, page 3277–3280, New York, NY, USA. Association for Computing Machinery.
- Paul A Kirschner, Simon J Buckingham-Shum, and Chad S Carr. 2012. *Visualizing argumentation: Software tools for collaborative and educational sense-making*. Springer Science & Business Media.
- Guillaume Klein, Yoon Kim, Yuntian Deng, Jean Senellart, and Alexander M. Rush. 2017. [Opennmt: Open-source toolkit for neural machine translation](#).
- Terry K Koo and Mae Y Li. 2016. A guideline of selecting and reporting intraclass correlation coefficients for reliability research. *Journal of chiropractic medicine*, 15(2):155–163.
- 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*, pages 7871–7880, Online. Association for Computational Linguistics.
- Manling Li, Lingyu Zhang, Heng Ji, and Richard J. Radke. 2019. [Keep meeting summaries on topic: Abstractive multi-modal meeting summarization](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 2190–2196, Florence, Italy. Association for Computational Linguistics.
- Chin-Yew Lin and Franz Josef Och. 2004. Automatic evaluation of machine translation quality using longest common subsequence and skip-bigram statistics. In *Proceedings of the 42nd Annual Meeting on Association for Computational Linguistics*, page 605. Association for Computational Linguistics.
- Chunyi Liu, Peng Wang, Jiang Xu, Zang Li, and Jieping Ye. 2019a. [Automatic dialogue summary generation for customer service](#). In *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining, KDD19*, page 1957–1965, New York, NY, USA. Association for Computing Machinery.
- H. Liu and P. Singh. 2004. [Conceptnet — a practical commonsense reasoning tool-kit](#). *BT Technology Journal*, 22(4):211–226.
- Yang Liu and Mirella Lapata. 2019. Text summarization with pretrained encoders. 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 3721–3731.
- Zhengyuan Liu, Angela Ng, Sheldon Lee, Ai Ti Aw, and Nancy F. Chen. 2019b. [Topic-aware pointer-generator networks for summarizing spoken conversations](#). *2019 IEEE Automatic Speech Recognition and Understanding Workshop (ASRU)*.
- Christopher Manning, Mihai Surdeanu, John Bauer, Jenny Finkel, Steven Bethard, and David McClosky. 2014. [The Stanford CoreNLP natural language processing toolkit](#). In *Proceedings of 52nd Annual Meeting of the Association for Computational Linguistics: System Demonstrations*, pages 55–60, Baltimore, Maryland. Association for Computational Linguistics.
- Joshua Maynez, Shashi Narayan, Bernd Bohnet, and Ryan McDonald. 2020. [On faithfulness and factuality in abstractive summarization](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 1906–1919, Online. Association for Computational Linguistics.
- K. McKeown, Lokesh Shrestha, and Owen Rambow. 2007. Using question-answer pairs in extractive summarization of email conversations. In *CICLing*.
- Amita Misra, Pranav Anand, Jean E. Fox Tree, and Marilyn Walker. 2015. [Using summarization to discover argument facets in online ideological dialog](#). In *Proceedings of the 2015 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 430–440, Denver, Colorado. Association for Computational Linguistics.
- Beverly Moser. 2001. An introduction to discourse analysis. theory and method.
- Gabriel Murray, S. Renals, and J. Carletta. 2005. Extractive summarization of meeting recordings. In *INTERSPEECH*.

- Gabriel Murray, Steve Renals, Jean Carletta, and Johanna Moore. 2006. [Incorporating speaker and discourse features into speech summarization](#). In *Proceedings of the Human Language Technology Conference of the NAACL, Main Conference*, pages 367–374, New York City, USA. Association for Computational Linguistics.
- Shashi Narayan, Shay B Cohen, and Mirella Lapata. 2018. Ranking sentences for extractive summarization with reinforcement learning. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pages 1747–1759.
- Romain Paulus, Caiming Xiong, and Richard Socher. 2018. [A deep reinforced model for abstractive summarization](#). In *International Conference on Learning Representations*.
- Ke Chen Qin, Lu Wang, and Joseph Kim. 2017. Joint modeling of content and discourse relations in dialogues. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 974–984.
- Alexander M. Rush, Sumit Chopra, and Jason Weston. 2015. [A neural attention model for abstractive sentence summarization](#). In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*, pages 379–389, Lisbon, Portugal. Association for Computational Linguistics.
- Harvey Sacks, Emanuel A Schegloff, and Gail Jefferson. 1978. A simplest systematics for the organization of turn taking for conversation. In *Studies in the organization of conversational interaction*, pages 7–55. Elsevier.
- Abigail See, Peter J. Liu, and Christopher D. Manning. 2017. [Get to the point: Summarization with pointer-generator networks](#). *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*.
- Guokan Shang, Wensi Ding, Zekun Zhang, Antoine Tixier, Polykarpos Meladianos, Michalis Vazirgiannis, and Jean-Pierre Lorré. 2018. [Unsupervised abstractive meeting summarization with multi-sentence compression and budgeted submodular maximization](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 664–674, Melbourne, Australia. Association for Computational Linguistics.
- Zhouxing Shi and Minlie Huang. 2019. A deep sequential model for discourse parsing on multi-party dialogues. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, pages 7007–7014.
- Kaiqiang Song, Chen Li, Xiaoyang Wang, Dong Yu, and Fei Liu. 2020. The ucf podcast summarization system at trec 2020. *arXiv preprint arXiv:2011.04132*.
- Matthew Stone, Una Stojnic, and Ernest Lepore. 2013. Situated utterances and discourse relations. In *Proceedings of the 10th International Conference on Computational Semantics (IWCS 2013)–Short Papers*, pages 390–396.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. In *Advances in neural information processing systems*, pages 5998–6008.
- Petar Veličković, Guillem Cucurull, Arantxa Casanova, Adriana Romero, Pietro Liò, and Yoshua Bengio. 2018. [Graph attention networks](#). In *International Conference on Learning Representations*.
- Lu Wang and Claire Cardie. 2013. Domain-independent abstract generation for focused meeting summarization. In *Proceedings of the 51st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1395–1405.
- Jiacheng Xu, Zhe Gan, Yu Cheng, and Jingjing Liu. 2020. [Discourse-aware neural extractive text summarization](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 5021–5031, Online. Association for Computational Linguistics.
- Jingqing Zhang, Yao Zhao, Mohammad Saleh, and Peter J Liu. 2019. Pegasus: Pre-training with extracted gap-sentences for abstractive summarization. *arXiv preprint arXiv:1912.08777*.
- Zhou Zhao, Haojie Pan, Changjie Fan, Yan Liu, Lintin Li, Min Yang, and Deng Cai. 2019. [Abstractive meeting summarization via hierarchical adaptive segmental network learning](#). In *The World Wide Web Conference, WWW '19*, page 3455–3461, New York, NY, USA. Association for Computing Machinery.
- Chenguang Zhu, William Hinthorn, Ruochen Xu, Qingkai Zeng, Michael Zeng, Xuedong Huang, and Meng Jiang. 2020a. Boosting factual correctness of abstractive summarization with knowledge graph. *ArXiv*, abs/2003.08612.
- Chenguang Zhu, Ruochen Xu, Michael Zeng, and Xuedong Huang. 2020b. A hierarchical network for abstractive meeting summarization with cross-domain pretraining. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, Online. Association for Computational Linguistics.

A Discourse Relation Distributions

We pre-trained a deep sequential model (Shi and Huang, 2019) on STAC Corpus (1,062 dialogues)

Model	ROUGE-1			ROUGE-2			ROUGE-L		
	F	P	R	F	P	R	F	P	R
BART (Lewis et al., 2020)	45.15	49.58	45.97	21.66	23.95	22.16	44.46	48.92	44.26
S-BART w. Discourse $\alpha = 0$ †	45.40	50.22	45.86	21.96	24.49	22.25	44.56	49.32	44.13
S-BART w. Action $\alpha = 0$ †	45.47	50.82	45.42	22.23	24.96	22.34	44.55	49.69	43.75
S-BART w. Discourse&Action $\alpha = 0$ †	45.59	51.47	45.09	22.42	25.51	22.27	44.67	50.24	43.52
S-BART w. Discourse $\alpha = 1$ †	45.89	51.34	45.87	22.50	25.26	22.33	44.83	49.93	44.17
S-BART w. Action $\alpha = 1$ †	45.67	50.25	46.44	22.39	24.70	22.96	44.86	49.29	44.75
S-BART w. Discourse&Action $\alpha = 1$ †	46.07	51.13	46.24	22.60	25.11	22.81	45.00	49.82	44.47

Table 8: Results on SAMSum Corpus. ROUGE-1, ROUGE-2 and ROUGE-L scores for different models on the test set. Results are averaged over three random runs. † means our methods.

Discourse Type	SAMSum	ADSC
Comment	19.3%	42.7%
Clarification	15.2%	13.3%
Elaboration	2.3%	0.1%
Acknowledgement	8.4%	0.9%
Explanation	2.8%	0.3%
Conditional	0.2%	0%
QA pair	21.5%	12.3%
Alternation	0.3%	0.6%
Result	5.5%	0.2%
Background	0.4%	0%
Narration	0.4%	0%
Correction	0.4%	1.1%
Continuation	0.9%	7.5%
Q-Elab	2.5%	0%
Parallel	0.9%	0%
Contrast	1.0%	19.5%

Table 9: The distribution of predicted discourse relation types on SAMSum Corpus and ADSC Corpus.

that with 0 under under all graph settings, suggesting utilizing more information from graphs would bring in more performance boosts.

(Asher et al., 2016) with default settings⁵ to get the link prediction and relation classification models to label discourse relations in SAMSum and ADSC corpus. The distribution of the relation types in two datasets were shown in Table 9. The major discourse relations in daily conversations are Comment, Clarification and QA pairs, while the main discourse relations in debate are Comment, Contrast, Clarification and QA pairs.

B Impact of Different ReZero Weight Initializations

We tested our structure-aware BART (*S-BART w. Discourse/Action*) within two ReZero settings: (i) initializing α from 0, (ii) initializing α from 1. And the results were shown in Table 8. *S-BART* with 1 as the initialized ReZero weight outperformed

⁵<https://github.com/shizhouxing/DialogueDiscourseParsing>