TOPIC-AWARE POINTER-GENERATOR NETWORKS FOR SUMMARIZING SPOKEN CONVERSATIONS

Zhengyuan Liu[†], Angela Ng[‡], Sheldon Lee[‡], Ai Ti Aw[†], Nancy F. Chen[†]*

[†]Institute for Infocomm Research, A*STAR, Singapore
[‡]Changi General Hospital, Singapore

ABSTRACT

Due to the lack of publicly available resources, conversation summarization has received far less attention than text summarization. As the purpose of conversations is to exchange information between at least two interlocutors, key information about a certain topic is often scattered and spanned across multiple utterances and turns from different speakers. This phenomenon is more pronounced during spoken conversations, where speech characteristics such as backchanneling and false-starts might interrupt the topical flow. Moreover, topic diffusion and (intra-utterance) topic drift are also more common in human-to-human conversations. Such linguistic characteristics of dialogue topics make sentence-level extractive summarization approaches used in spoken documents ill-suited for summarizing conversations. Pointer-generator networks have effectively demonstrated its strength at integrating extractive and abstractive capabilities through neural modeling in text summarization. To the best of our knowledge, to date no one has adopted it for summarizing conversations. In this work, we propose a topic-aware architecture to exploit the inherent hierarchical structure in conversations to further adapt the pointer-generator model. Our approach significantly outperforms competitive baselines, achieves more efficient learning outcomes, and attains more robust performance.

Index Terms— Dialogue, summarization, neural networks, attention mechanism, conversation technology

1. INTRODUCTION

Automatic summarization condenses lengthy materials into shorter versions, which focuses on the essential information and the overall meaning. Such summaries can enable users to browse and digest information more effectively and efficiently, which is especially useful in the digital era of information overload. Thus, summarization has attracted much research attention over the years. When humans generate summaries, they usually first comprehend the entire content, and then extract and concatenate the keywords or salient sentences. To obtain more succinct and readable results, they further distill and rephrase the most important information [1]. This gives rise to the two main paradigms in automatic summarization: extractive and abstractive. Recently, with sophisticated neural architectures, representation learning of linguistic elements and large-scale available corpora, data-driven approaches have made much progress in both two paradigms [2, 3, 4]. Most work in text summarization focuses on documents such as news or academic articles [5, 6]. On the contrary, due to the scarcity of publicly available corpora with ground-truth summaries, the speech summarization task of human-to-human spoken conversations has received far less attention even though there is high industry demand and many potential applications across different domains [7, 8].

Different from passages, human-to-human spoken conversations are a dynamic flow of information exchange, which is often informal, verbose and repetitive, sprinkled with false-starts, backchanneling, reconfirmations, hesitations, and speaker interruptions [9]. The key information about a certain topic is often at the sub-sentence or sub-utterance level, and scattered and spanned across multiple utterances and turns from different speakers, leading to lower information density, more diffuse topic coverage, and (intra-utterance) topic drifts. These spoken characteristics pose technical challenges for sentence-level extractive approaches that would inevitably include unnecessary spans of words in the generated summaries. Pointer-generator networks [10], a neural sequence-to-sequence design [11], produce summaries via word-level extraction and abstractive generation. Thus, we propose to exploit its advantages to tackle the aforementioned challenges. Meanwhile, although conversations are often less structured than passages, they are inherently organized around the dialogue topics in a coarse-grained structure [9]. Topic segmentation has also been shown to be useful in dialogue-based information retrieval [12]. Such prior analyses and investigations inspire us to augment pointer-generator networks with a topic-level attention mechanism to more ele-

gantly attend to the underlying (yet often interrupted) topical flow in human-to-human conversations.

The dialogue setting in this work is based on a real-world scenario where nurses discuss symptom information with post-discharge patients to follow-up with their status. The conversation summarization task here is targeted at automatically generating notes that describe the symptoms the patients are experiencing. Our proposed topic-aware pointer-generator framework exploits the inherent hierarchical structure in dialogues and is able to address the technical challenges of spoken dialogue summarization: empirical results outperform competitive baseline architectures significantly, while achieving more efficient and robust learning outcomes.

2. RELATED WORK

In speech and text summarization, traditional approaches are studied more on the extractive methods, utilizing rule-based, statistical [13] and graph-based [14] algorithms with various linguistic features like lexical similarity [15], semantic structure [7], or discourse relation [16]. Recently, end-to-end neural approaches have been widely adopted in text summarization due to their capability, flexibility and scalability. Neural extractive models adopt sentence labeling or ranking strategies [17, 18], semantic vector representations [19] and sequential context modeling [20]. For abstractive tasks, sequence-to-sequence models use a neural decoder to generate informative and readable summaries word-by-word [3]. Various methods have been proposed for improvement: the attention mechanism helps the decoder concentrate on appropriate parts of source content [21]; the pointer mechanism is effective in handling out-of-vocabulary words [10]; the coverage mechanism is used to reduce generative repetitions [22]. Extractive-abstractive models have also been proposed to obtain better results [23, 24].

While text summarization focuses on passages such as news articles [5, 25] and academic publications [6], speech summarization has been investigated in monologues (such as broadcasts [26] and lectures [7]) and multi-party dialogues such as meetings [27]. Recently, neural modeling approaches have also been adopted: Goo *et al.* [8] used a sequence-to-sequence model to write headlines for meetings, and Liu *et al.* [4] used a hierarchical extractive model to summarize monologue spoken documents. In this paper, we propose a topic-aware pointer-generator architecture for hierarchical context modeling to generate summary notes of spoken conversations.

3. CONVERSATION CORPUS AND SETUP

We sampled 100k dialogues as the training set and another distinct 1k dialogues as the validation set according to Section 3.1. While the training and validation sets were constructed using simulated data, the test set was derived from 490 multi-turn dialogues that took place between nurses and patients in the healthcare setting. Topic segmentation (Section 3.2) and ground-truth summary construction (Section 3.3) were conducted on all these subsets.

3.1. Nurse-to-Patient Dialogue Data

This corpus was inspired by a pilot set of conversations that took place in the clinical setting where nurses inquire about symptoms of patients [28]. Linguistic structures at the semantic, syntactic, discourse and pragmatic levels were systemically abstracted from these conversations to construct templates for automatically simulating multi-turn dialogues [29]. The informal and spontaneous styles of spoken interactions such as interlocutor interruption, backchanneling, hesitation, false-starts, repetition, and topic drift were preserved (see Figure 1a for an example).

A team of linguistically trained personnel refined, substantiated, and corrected the simulated dialogues by enriching verbal expressions through considering paraphrasing, different regional English speaking styles (American, British, and Asian) through word usage and sentence patterns, validating logical correctness through considering if the dialogues were natural and not disobeying common sense, and verifying the clinical content by consulting certified nurses. These conversations cover 9 topics/symptoms (e.g., headache, cough).
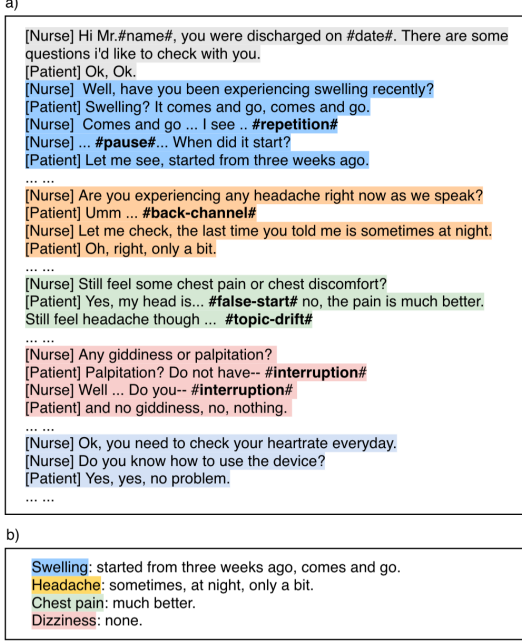
3.2. Topic Segmentation

In dialogue analysis, a change in topic corresponds to a change in cognitive attention acknowledged and acted upon by speakers, which is usually related to content themes [12]. In this work, we specify the dialogue topics according to the symptoms in nurse-to-patient conversations. Figure 1a shows an example of different topic segments, where each spans across different utterances and speakers. Note that various types of spoken characteristics (e.g. false start) could break up topical congruence.

A rule-based lexical algorithm was used to detect the boundaries between dialogue topics. Labels $\langle s \rangle$ and $\langle /s \rangle$ were respectively added before and after each topic segment. Position indices of segment labels are used in Section 4.3. Human verification was conducted to ensure quality control.

3.3. Ground-Truth Summaries

The goal of this conversation summarization task is to obtain a concise description characterizing different attributes of a specified symptom. For this particular clinical scenario, the summary notes are preferred to be represented in a very structured format to facilitate indexing, searching, retrieving and extracting in a variety of downstream workflow applications (e.g., decision support, medical triage). Thus, paraphrases of a particular symptom are represented using the same entity, e.g., *"shortness of breath"* and *"breathlessness"* are both represented as symptom *"breathlessness"*.

a)

[Nurse] Hi Mr.#name#, you were discharged on #date#. There are some questions i'd like to check with you.
[Patient] Ok, Ok.
[Nurse] Well, have you been experiencing swelling recently?
[Patient] Swelling? It comes and go, comes and go.
[Nurse] Comes and go ... I see .. **#repetition#**
[Nurse] ... **#pause#**... When did it start?
[Patient] Let me see, started from three weeks ago.
... ...
[Nurse] Are you experiencing any headache right now as we speak?
[Patient] Umm ... **#back-channel#**
[Nurse] Let me check, the last time you told me is sometimes at night.
[Patient] Oh, right, only a bit.
... ...
[Nurse] Still feel some chest pain or chest discomfort?
[Patient] Yes, my head is... **#false-start#** no, the pain is much better. Still feel headache though ... **#topic-drift#**
... ...
[Nurse] Any giddiness or palpitation?
[Patient] Palpitation? Do not have-- **#interruption#**
[Nurse] Well ... Do you-- **#interruption#**
[Patient] and no giddiness, no, nothing.
... ...
[Nurse] Ok, you need to check your heartrate everyday.
[Nurse] Do you know how to use the device?
[Patient] Yes, yes, no problem.
... ...

b)

Swelling: started from three weeks ago, comes and go.
Headache: sometimes, at night, only a bit.
Chest pain: much better.
Dizziness: none.

Fig. 1. a) A dialogue example of multi-turn conversation on nurse-patient symptom discussion. Colored spans are utterances of the same topic. Spoken characteristics are preserved and represented in bold font; b) Generated ground-truth summary. Colored spans indicate corresponding topics in the given dialogue example.

The summary format is shown in Figure 1b, where each symptom is listed separately with corresponding attributes such as frequency of symptom or severity of symptom that were mentioned in the conversations. If a symptom was mentioned, it will be included in the summary. If there is no signal of a symptom (e.g., *"cough"*) in the discussion between the nurse and the patient, the summary for the symptom is represented as *"Cough: none"*, while the others would be recorded with key information from the dialogue, e.g., *"Headache: every night, only a bit"*. Human verification was conducted to ensure quality control.

4. APPROACH

In a sequence-to-sequence (seq2seq) model, the encoder receives a token sequence of content $x = \{x_0, x_1, ..., x_n\}$ of length n, and the decoder outputs a token sequence of summary $y = \{y_0, y_1, ..., y_m\}$ of length m. The decoder generates words step by step, and previously generated words are encoded to provide contextual information at each time step. The task is to learn a function f with a parameter set θ that maximizes the probability to generate readable and meaningful output text. In this section, we (1) describe two baselines: an attentive seq2seq model and a pointer-generator network, and (2) demonstrate how we integrate topic-level attention neural mechanisms.

4.1. Attentive Seq2Seq Architecture

The attentive seq2seq model is similar to that in [21]. It adds an attention layer to a vanilla seq2seq network, to filter out unnecessary contextual information at each decoding step.

Sequence Encoding: Given a document input x (one-hot word representation), an embedding layer converts it to vector representations by a look-up operation using the embedding matrix, obtaining a vector sequence $v = \{v_0, v_1, ..., v_n\}$, where $v_i \in R^d$ and d is the embedding dimension. Then, a bi-directional long short-term memory (Bi-LSTM) layer [30] is used to encode v_i with forward and backward temporal dependencies, as $\overrightarrow{h_i}$ and $\overleftarrow{h_i}$, respectively, and concatenate them as the hidden representation $h = \{h_0, h_1, ..., h_n\} \in R^{n \times 2d'}$, where d' is the hidden dimension size.

$$\overrightarrow{h_i} = \overrightarrow{\text{LSTM}}(v_i, \overrightarrow{h_{i-1}}); \quad \overleftarrow{h_i} = \overleftarrow{\text{LSTM}}(v_i, \overleftarrow{h_{i+1}}) \quad (1)$$

$$h_i = [\overrightarrow{h_i}; \overleftarrow{h_i}] \quad (2)$$

Sequence Decoding with Attention: The decoder is a single-layer unidirectional LSTM, generating words step by step. For each decoding step t, it receives the word embedding of the previous token y_{t-1}, then calculates the decoder state s_t. Attention scoring is conducted through concatenation [31]:

$$a^t = \text{softmax}(W_e \tanh(W_{attn}[h_i; s_t] + b_{attn})) \quad (3)$$

where W_e, W_{attn} and b_{attn} are trainable parameters, and $[;]$ is the concatenation operation. Attention scores can be viewed as the importance over the input content, guiding the decoder to concentrate on the appropriate positions of context for generating the next word. Next, these attention scores are used to produce a weighted-pooling of the encoded hidden states, as a fixed-size representation of what has been read from the source for this step, namely the context vector h_t^c.

Then, h_t^c is concatenated with the decoder state s_t and fed through two dense layers to produce a distribution on the vocabulary:

$$p_{vocab}^t = \text{softmax}(W''(W'[s_t; h_t^c] + b') + b'') \quad (4)$$

where W', W'', b' and b'' are trainable parameters. p_{vocab}^t is a probability distribution over all words in the vocabulary, which will be used to generate decoded tokens.

4.2. Pointer-Generator Networks

Pointer-generator networks are a variant of seq2seq architecture by adding a pointer network [32]. Aside from generating words over a fixed vocabulary, the model is able to copy words via pointing to the source content, thus bypassing out-of-vocabulary issues. In the pointer-generator model, the sequence encoding representation h, attention distribution a^t and context vector h_t^c are calculated as in Section 4.1. We describe how the pointer-generator model achieves word-level extraction and abstractive generation below.
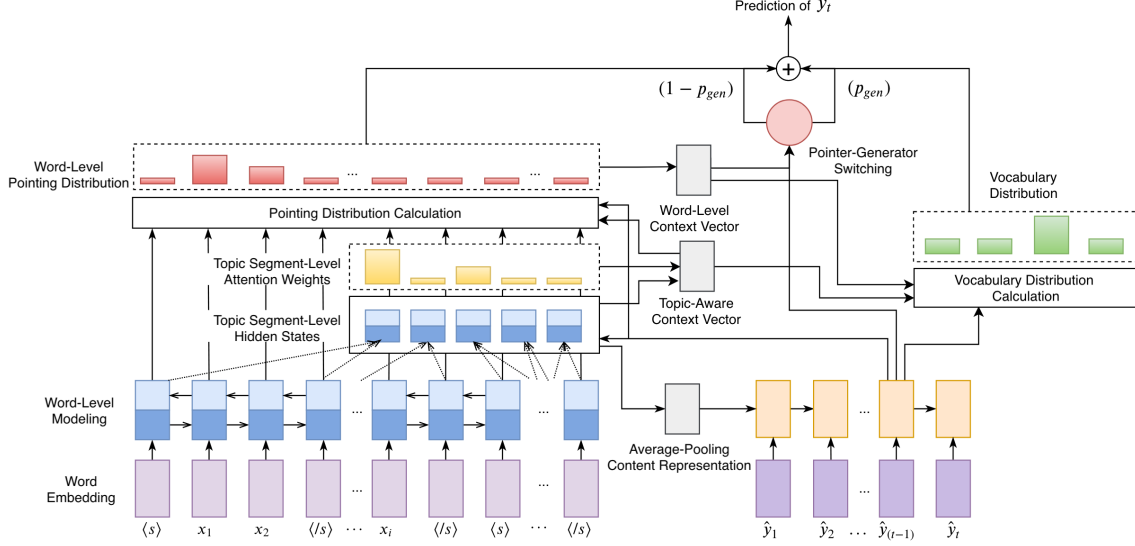
Fig. 2. The architecture of our proposed topic-aware pointer-generator network.

Pointer Mechanism: To directly copy words from the source content, the attentive distribution a^t in Equation (3) is regarded as the copy probability on input sequence x, specifically, the token on position with max probability will be extracted as the output.

Pointer-Generator Switching: At each decoding step, there is a switching probability to determine whether to generate a token from the fixed vocabulary or copy one from the source. $p_{gen} \in [0, 1]$ for step t is calculated from the context vector h_c^t, the decoder input y_{t-1} and the decoder state s_t:

$$p_{gen}^t = \sigma(W_{ptr}[h_c^t; y_{t-1}; s_t] + b_{ptr}) \qquad (5)$$

where W_{ptr} and b_{ptr} are trainable parameters and σ is the sigmoid function. Then, p_{gen}^t is used as a soft switch to choose between copying a word from the input sequence with attention distribution a^t, or generating a word from the vocabulary distribution p_{vocab}^t in Equation (4). For each sample, we extend the vocabulary with the unique words from the input content. We obtain the following probability distribution over the extended vocabulary:

$$p_{output}^t = p_{gen}^t \times p_{vocab}^t(w) + (1 - p_{gen}^t) \times \sum_{i:w_i=w} a_i^t \quad (6)$$

Contrary to the vanilla seq2seq model, which is restricted to a fixed vocabulary, the ability to copy and generate words is one of the primary advantages of the pointer-generator design. Another improvement over vanilla seq2seq is its coverage mechanism that avoids repetitions during decoding. More details can be found in [10].

4.3. Topic-Aware Attentive Architecture

To exploit topical structure in dialogue modeling with a hierarchical architecture, another attention layer is introduced to score topic-level attention to obtain a topic-aware context representation. We delineate how we integrate topic-aware attention to the baseline models below (see Figure 2).

Topic-Level States: We obtain the representations h^{seg} of topic segments by collecting hidden states from h in equation (1) with the topic-level segment position indices $t^{seg} = \{t_0^{seg}, t_1^{seg}, ...t_k^{seg}\}$, where k is the topic segment number of the dialogue content. For the output of bi-directional LSTM, we collect the states of forward and backward directions and then concatenate them into one.

Initial Decode State: In the two baseline models, the last hidden state of sequence encoding representation is used as the initial state s_0 fed to the decoder. Here, we denote average-pooling of topic-level states h^{seg} as s_0.

Topic-Aware Contextual Representation: In each decoding time step t, we calculate topic-level attention and use topic-aware context vectors as the guide for the fine-grained word-level prediction. First, the topic-level attention scores are calculated via dense layers and softmax normalization:

$$a_t^{seg} = \text{softmax}(W_e^{seg}\tanh(h_i W_{seg} s_t + b_{seg})) \qquad (7)$$

Next, we multiply the attention score with topic-level states to obtain the topic-aware context vector:

$$h_t^* = \sum_i^k h^{seg} a_t^{seg} \qquad (8)$$

Then, the pointing distribution in Equation (3) and vocabulary distribution in Equation (4) are influenced by the topic-aware contextual representation, and final output is produced as in Equation (6):

$$a^t = \text{softmax}(W_e\tanh(W_{attn}[h_i; h_t^*; s_t] + b_{attn})) \qquad (9)$$

$$p_{vocab}^t = \text{softmax}(W''(W'[s_t; h_t^c; h_t^*] + b') + b'') \qquad (10)$$

Model	ROUGE-1			ROUGE-2			ROUGE-L		
	F1	Precision	Recall	F1	Precision	Recall	F1	Precision	Recall
Attn Seq2Seq	0.4494	0.3495	0.8312	0.3444	0.2664	0.6582	0.4074	0.3171	0.7542
Attn Seq2Seq + TA	0.5041	0.3922	0.8230	0.3874	0.2998	0.6491	0.4560	0.3595	0.7482
PG-Net	0.5345	0.4238	0.8653	0.4311	0.3341	0.7220	0.4949	0.3843	0.8159
PG-Net + TA (Proposed)	**0.5862**	0.4803	0.8703	**0.4906**	0.3928	0.7550	**0.5496**	0.4423	0.8260

Table 1. Evaluation results of baselines and the proposed model.

5. EXPERIMENTS

5.1. Training Setup

The experiments are conducted on the nurse-to-patient conversation corpus described in Section 3. We implemented an attentive seq2seq (*Attn Seq2Seq*) and a pointer-generator network (*PG-Net*) as baselines, a topic-aware attentive seq2seq (*Attn Seq2Seq+TA*) model as control, and our proposed topic-aware pointer-generator model (*PG-Net+TA*).

Cross entropy is used to measure the loss between prediction and ground-truth. For time step t, the negative log likelihood of the target word $\hat{y}_t$ is defined as:

$$loss_t = -\log P_\theta(\hat{y}_t|x) \qquad (11)$$

and the overall loss is the sum from all the time steps. Teacher forcing strategy is applied: during training, the input is the previous word from the ground truth; at test time, the input is the previous word predicted by the decoder.

In our setup, for both the encoder and decoder, the dimension of word embeddings and hidden states was set to 200. We adopted pre-trained word embedding Glove [19], and out-of-vocabulary words and segment labels were initiated with random vectors. Embedding weight sharing strategy was applied by sharing the same embedding matrix W_{emb} for both encoder and decoder. This sharing significantly reduced parameter size and boosted the performance by reusing the semantic and syntactic information in one embedding space [22]. The learning rate was fixed to 0.0001 and the batch size was set to 32. We adopted gradient clipping with a maximum gradient norm of 2.0. Adam algorithm [33] was used for stochastic optimization, with $\beta_1 = 0.9, \beta_2 = 0.99$. The vocabulary size was $10K$. We limited source contents to 300 tokens and the decoding length to 100 tokens. We adopted early-stop strategy with validation in each training epoch. During testing, we set the beam search size to 5.

5.2. Empirical Evaluations

5.2.1. Evaluation I: Effectiveness

Summarization performance is measured using ROUGE-1, ROUGE-2 and ROUGE-L scores [34], as shown in Table 1. Between the two baselines, the pointer-generator model obtains higher performance than attentive seq2seq model. Qual-
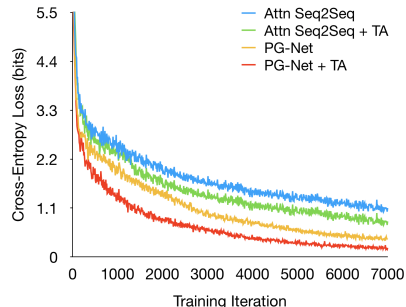

Fig. 3. Evaluation results of learning efficiency.

itative analysis (Section 5.3.1) shows that a certain proportion of generated tokens are directly copied from the source content, demonstrating the effectiveness of using the pointer mechanism in our task. Moreover, with topic-aware attentive modeling, both baseline models obtain significant improvement, and the proposed *PG-Net+TA* achieves the best performance. The gains are more prominent for the precision scores (see Table 1), indicating that the proposed approach generates fewer unnecessary tokens while preserving key information in the generated summary.

5.2.2. Evaluation II: Learning Efficiency

To evaluate the learning efficiency of the models, we recorded their loss values during training. As shown in Figure 3, in the first 7000 batches of iterations, loss of the pointer-generator decreases faster than that of attentive seq2seq. Moreover, by adding topic-level attention, both attentive seq2seq and pointer-generator are improved, and our proposed *PG-Net+TA* performs best.

Having demonstrated the strength of pointer-generator networks over attentive seq2seq model, we will focus on the former in the following experiments.

5.2.3. Evaluation III: Performance/Model Robustness

Spoken conversations are often verbose with low information density scattered with topics not central to the main dialogue theme, especially since speakers chit-chat and get distracted during task-oriented discussions. To evaluate such scenarios, we adopted model-independent ADDSENT [35], where we randomly extracted sentences from SQuAD and inserted

Fig. 4. a) Switching between pointer & generator: Tokens with higher switching probability are darker in color and are generated from the vocabulary. b) Visualizing topic-level attention scores. The darker the shade, the higher the scores.

Model	ROUGE-1	ROUGE-2	ROUGE-L
PG-Net	42.33	28.23	38.14
	(-11.12)	(-14.88)	(-11.35)
PG-Net + TA	49.70	39.32	45.37
	(-8.92)	(-9.74)	(-9.59)

Table 2. F1 scores from lengthy-sample evaluation. Bracketed values denote absolute decrease of model performance in Section 5.2.1.

them before or after topically coherent segments. The average length of the augmented test set, is increased from 300 to 900. As shown in Table 2, topic-level attention helps the pointer-generator model become more robust to lengthy samples.

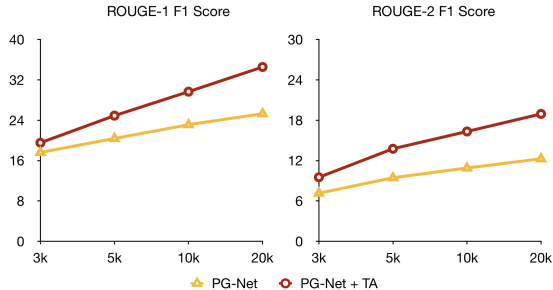### 5.2.4. *Evaluation IV: Low Resource Training*

Limited amount of training data is a major pain point for dialogue-based tasks, as it is time-consuming and labor-intensive to collect and annotate natural dialogues at a large-scale. We expect our model to perform better in low resource scenarios because it can take advantage of the inherently hierarchical dialogue structure with induction bias. We conducted experiments over a range of training sizes (from 3k to 20k). As shown in Figure 5, our proposed *PG-Net+TA* lead to steeper learning curves and always outperforms the baseline.

5.3. Visualization Analysis

In this section, we probe deeper into the proposed neural architecture to examine the innerworkings of (1) how the pointer and generator switches, and (2) how topic-level attention interacts with the decoded sequence.

5.3.1. *Pointer-Generator Switching*

We illustrate the probability of the pointer-generator switching p_{gen} in Equation (5) that indicates the probability of words generated from the vocabulary to show how the proposed model summarizes dialogue. One summary example produced from a dialogue with three topics is shown in Figure 4a: the attribute information of a symptom and segment tags are directly copied from the source content, while symptom

Fig. 5. Evaluation results of low resource training.

entities and punctuation tokens are generated from the vocabulary lexicon. This generation behavior resonates with the rationale used in constructing the ground-truth summary in Section 3.3, enabling the model to normalize symptom entities and handle out-of-vocabulary words in symptom attributes.

5.3.2. *Topic-Level Attention Scoring*

To show how the proposed framework conducts topic-aware contextual modeling, we illustrate topic-level attention scores a^{seg} in Equation (7). As shown in Figure 4b, during the summary decoding process of a dialogue with three topics, at each step, the model concentrates on one topic segment. We also observe smooth topic transition from the attention layer, which aligns well with the topical flow of the dialogue content. Such topic level modeling can help improve summarization performance by filtering out nonessential details at the word-level modeling layers.

6. CONCLUSION

In this work, we automatically summarized spoken dialogues from nurse-to-patient conversations. We presented an effective and efficient neural architecture that integrates topic-level attention mechanism in pointer-generator networks, utilizing the hierarchical structure of dialogues. We demonstrated that the proposed model significantly outperforms competitive baselines, obtains more efficient learning outcomes, is robust to lengthy dialogue samples, and performs well when there is limited training data.

7. REFERENCES

[1] Hongyan Jing and Kathleen R. McKeown, "The decomposition of human-written summary sentences," in *Proceedings of the 22Nd Annual International ACM SIGIR Conference on Research and Development in Information Retrieval*, New York, NY, USA, 1999, SIGIR '99, pp. 129–136, ACM.

[2] Ramesh Nallapati, Feifei Zhai, and Bowen Zhou, "Summarunner: A recurrent neural network based sequence model for extractive summarization of documents," in *Thirty-First AAAI Conference on Artificial Intelligence*, 2017.

[3] Ramesh Nallapati, Bowen Zhou, Cicero dos Santos, Çağlar Gulçehre, and Bing Xiang, "Abstractive text summarization using sequence-to-sequence RNNs and beyond," in *Proceedings of The 20th SIGNLL Conference on Computational Natural Language Learning*, Berlin, Germany, Aug. 2016, pp. 280–290, Association for Computational Linguistics.

[4] T. Liu, S. Liu, and B. Chen, "A hierarchical neural summarization framework for spoken documents," in *ICASSP 2019 - 2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, May 2019, pp. 7185–7189.

[5] Karl Moritz Hermann, Tomas Kocisky, Edward Grefenstette, Lasse Espeholt, Will Kay, Mustafa Suleyman, and Phil Blunsom, "Teaching machines to read and comprehend," in *Advances in neural information processing systems*, 2015, pp. 1693–1701.

[6] Noemie Elhadad, M-Y Kan, Judith L Klavans, and Kathleen R McKeown, "Customization in a unified framework for summarizing medical literature," *Artificial intelligence in medicine*, vol. 33, no. 2, pp. 179–198, 2005.

[7] Justin Jian Zhang, Ho Yin Chan, and Pascale Fung, "Improving lecture speech summarization using rhetorical information," in *2007 IEEE Workshop on Automatic Speech Recognition & Understanding (ASRU)*. IEEE, 2007, pp. 195–200.

[8] Chih-Wen Goo and Yun-Nung Chen, "Abstractive dialogue summarization with sentence-gated modeling optimized by dialogue acts," in *2018 IEEE Spoken Language Technology Workshop (SLT)*. IEEE, 2018, pp. 735–742.

[9] Harvey Sacks, Emanuel A Schegloff, and Gail Jefferson, "A simplest systematics for the organization of turn taking for conversation," in *Studies in the organization of conversational interaction*, pp. 7–55. Elsevier, 1978.

[10] Abigail See, Peter J. Liu, and Christopher D. Manning, "Get to the point: Summarization with pointer-generator networks," in *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, Vancouver, Canada, July 2017, pp. 1073–1083, Association for Computational Linguistics.

[11] Ilya Sutskever, Oriol Vinyals, and Quoc V Le, "Sequence to sequence learning with neural networks," in *Advances in neural information processing systems*, 2014, pp. 3104–3112.

[12] Narjès Boufaden, Guy Lapalme, and Yoshua Bengio, "Topic segmentation: A first stage to dialog-based information extraction," in *In Natural Language Processing Pacific Rim Symposium, NLPRS01*. Citeseer, 2001.

[13] Kai Hong and Ani Nenkova, "Improving the estimation of word importance for news multi-document summarization," in *Proceedings of the 14th Conference of the European Chapter of the Association for Computational Linguistics*, 2014, pp. 712–721.

[14] Günes Erkan and Dragomir R. Radev, "Lexrank: Graph-based lexical centrality as salience in text summarization," *J. Artif. Int. Res.*, vol. 22, no. 1, pp. 457–479, Dec. 2004.

[15] Horacio Saggion and Thierry Poibeau, "Automatic text summarization: Past, present and future," in *Multi-source, multilingual information extraction and summarization*, pp. 3–21. Springer, 2013.

[16] T. Hirao, M. Nishino, Y. Yoshida, J. Suzuki, N. Yasuda, and M. Nagata, "Summarizing a document by trimming the discourse tree," *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 23, no. 11, pp. 2081–2092, Nov 2015.

[17] Chris Kedzie, Kathleen McKeown, and Hal Daumé III, "Content selection in deep learning models of summarization," in *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, Brussels, Belgium, Oct.-Nov. 2018, pp. 1818–1828, Association for Computational Linguistics.

[18] Shashi Narayan, Shay B. Cohen, and Mirella Lapata, "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)*, New Orleans, Louisiana, June 2018, pp. 1747–1759, Association for Computational Linguistics.

[19] Jeffrey Pennington, Richard Socher, and Christopher Manning, "Glove: Global vectors for word representation," in *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, Doha, Qatar, Oct. 2014, pp. 1532–1543, Association for Computational Linguistics.

[20] Sepp Hochreiter and Jürgen Schmidhuber, "Long short-term memory," *Neural computation*, vol. 9, no. 8, pp. 1735–1780, 1997.

[21] Sumit Chopra, Michael Auli, and Alexander M. Rush, "Abstractive sentence summarization with attentive recurrent neural networks," in *Proceedings of the 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, San Diego, California, June 2016, pp. 93–98, Association for Computational Linguistics.

[22] Romain Paulus, Caiming Xiong, and Richard Socher, "A deep reinforced model for abstractive summarization," in *Proceedings of the 6th International Conference on Learning Representations*, 2018.

[23] Wan-Ting Hsu, Chieh-Kai Lin, Ming-Ying Lee, Kerui Min, Jing Tang, and Min Sun, "A unified model for extractive and abstractive summarization using inconsistency loss," in *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, Melbourne, Australia, July 2018, pp. 132–141, Association for Computational Linguistics.

[24] Sebastian Gehrmann, Yuntian Deng, and Alexander Rush, "Bottom-up abstractive summarization," in *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, Brussels, Belgium, Oct.-Nov. 2018, pp. 4098–4109, Association for Computational Linguistics.

[25] Alexander M. Rush, Sumit Chopra, and Jason Weston, "A neural attention model for abstractive sentence summarization," in *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*, Lisbon, Portugal, Sept. 2015, pp. 379–389, Association for Computational Linguistics.

[26] Sameer Maskey and Julia Hirschberg, "Comparing lexical, acoustic/prosodic, structural and discourse features for speech summarization," in *Ninth European Conference on Speech Communication and Technology*, 2005.

[27] Shasha Xie and Yang Liu, "Improving supervised learning for meeting summarization using sampling and regression," *Computer Speech & Language*, vol. 24, no. 3, pp. 495–514, 2010.

[28] Zhengyuan Liu, Hazel Lim, Nur Farah Ain Suhaimi, Shao Chuen Tong, Sharon Ong, Angela Ng, Sheldon Lee, Michael R. Macdonald, Savitha Ramasamy, Pavitra Krishnaswamy, Wai Leng Chow, and Nancy F. Chen, "Fast prototyping a dialogue comprehension system for nurse-patient conversations on symptom monitoring," in *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, Minneapolis, June 2019, pp. 24–31, Association for Computational Linguistics.

[29] Pararth Shah, Dilek Hakkani-Tür, Bing Liu, and Gokhan Tür, "Bootstrapping a neural conversational agent with dialogue self-play, crowdsourcing and on-line reinforcement learning," in *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 3 (Industry Papers)*, New Orleans - Louisiana, June 2018, pp. 41–51, Association for Computational Linguistics.

[30] Mike Schuster and Kuldip K Paliwal, "Bidirectional recurrent neural networks," *IEEE Transactions on Signal Processing*, vol. 45, no. 11, pp. 2673–2681, 1997.

[31] Thang Luong, Hieu Pham, and Christopher D. Manning, "Effective approaches to attention-based neural machine translation," in *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*, Lisbon, Portugal, Sept. 2015, pp. 1412–1421, Association for Computational Linguistics.

[32] Oriol Vinyals, Meire Fortunato, and Navdeep Jaitly, "Pointer networks," in *Advances in Neural Information Processing Systems 28*, C. Cortes, N. D. Lawrence, D. D. Lee, M. Sugiyama, and R. Garnett, Eds., pp. 2692–2700. Curran Associates, Inc., 2015.

[33] Diederik P Kingma and Jimmy Ba, "Adam: A method for stochastic optimization," in *Proceedings of the 3rd International Conference for Learning Representations*, 2015.

[34] Chin-Yew Lin, "ROUGE: A package for automatic evaluation of summaries," in *Text Summarization Branches Out*, Barcelona, Spain, July 2004, pp. 74–81, Association for Computational Linguistics.

[35] Robin Jia and Percy Liang, "Adversarial examples for evaluating reading comprehension systems," in *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, Copenhagen, Denmark, Sept. 2017, pp. 2021–2031, Association for Computational Linguistics.