

Word Graph Guided Summarization for Radiology Findings

Jinpeng Hu^{♣,♡}, Jianling Li[◇], Zhihong Chen^{♣,♡}, Yaling Shen[♣],
Yan Song^{♣,♡}, Xiang Wan[♡], Tsung-Hui Chang^{♣,♡}

[♣]The Chinese University of Hong Kong (Shenzhen),

[♡]Shenzhen Research Institute of Big Data, [◇]National University of Defense Technology

[♣]{jinpenghu, zhihongchen, yalingshen}@link.cuhk.edu.cn

[◇]jianlingl@nudt.edu.cn [♡]wanxiang@sribd.cn

[♣]{songyan, changtsunghui}@cuhk.edu.cn

Abstract

Radiology reports play a critical role in communicating medical findings to physicians. In each report, the impression section summarizes essential radiology findings. In clinical practice, writing impression is highly demanded yet time-consuming and prone to errors for radiologists. Therefore, automatic impression generation has emerged as an attractive research direction to facilitate such clinical practice. Existing studies mainly focused on introducing salient word information to the general text summarization framework to guide the selection of the key content in radiology findings. However, for this task, a model needs not only capture the important words in findings but also accurately describe their relations so as to generate high-quality impressions. In this paper, we propose a novel method for automatic impression generation, where a word graph is constructed from the findings to record the critical words and their relations, then a **Word Graph guided Summarization model (WGSUM)** is designed to generate impressions with the help of the word graph. Experimental results on two datasets, OPENI and MIMIC-CXR, confirm the validity and effectiveness of our proposed approach, where the state-of-the-art results are achieved on both datasets. Further experiments are also conducted to analyze the impact of different graph designs to the performance of our method.¹

1 Introduction

A radiology report usually contains a findings section (FINDINGS) describing detailed medical observations and an impression section (IMPRESSION) summarizing the most critical observations. In practice, the IMPRESSION is an essential part and

¹Our code and the best performing models are released at <https://github.com/cuhksz-nlp/WGSUM>.

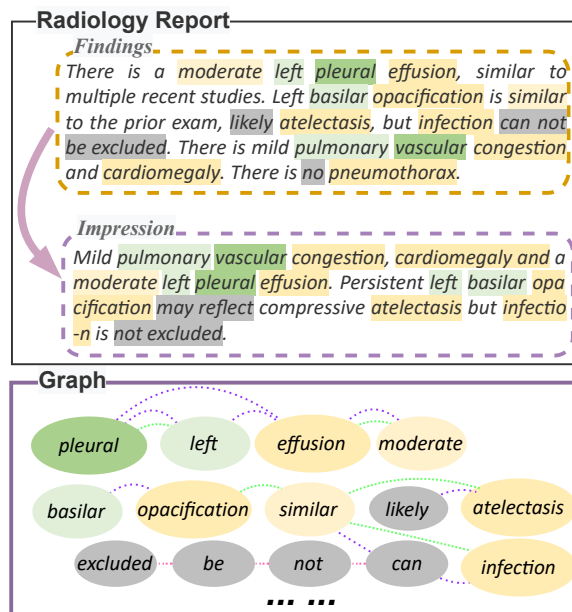


Figure 1: An example radiology report including its FINDINGS and IMPRESSION, and a word graph constructed from the FINDINGS. Different colored edges in the graph represent different types of relations. The curved arrow indicates the AIG task to generate the IMPRESSION from the FINDINGS.

plays an important role in delivering critical findings to clinicians. Therefore, summarizing FINDINGS helps to locate the most prominent observations so that the automatic process of doing so greatly eases the workload of radiologists. Recently, many methods are proposed for automatic impression generation (AIG) (Hassanpour and Langlotz, 2016; Zhang et al., 2018; Gharebagh et al., 2020), which are mainly based on the sequence-to-sequence architecture with specific designs for the characteristics of this task. For example, MacAvaney et al. (2019) employed clinical terms within the FINDINGS as key information to enhance AIG. Based on this work, Gharebagh et al. (2020) further proposed to identify the importance of these clinical terms and selected the most salient ones to facilitate the recognition of significant content

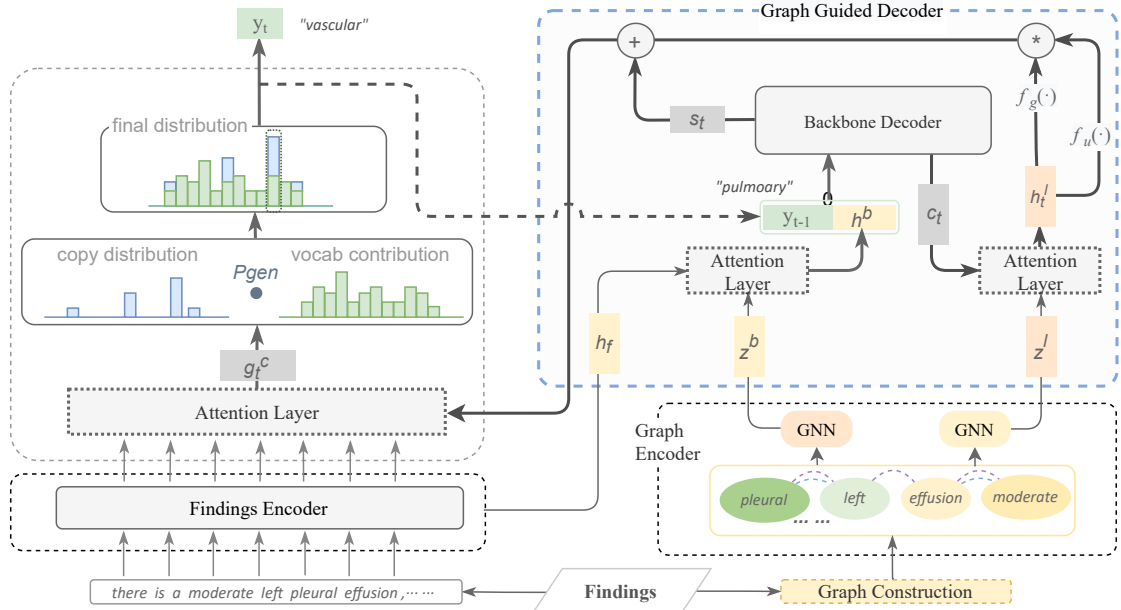


Figure 2: The overall architecture of our proposed method with an example input and sample outputs at the $t - 1$ and t step. The FINDINGS Encoder and the Graph Encoder are shown in different grey dashed boxes with their details omitted. The Graph Guided Decoder is illustrated in the blue dashed box.

so as to improve the performance. Although these efforts are able to find the important words to promote AIG, less attention is paid to the important aspect on leveraging the relation information among them. For example, in Figure 1, the observation word “*effusion*” and its modifier word “*moderate*” have a relation (which describes the severity of the symptoms) in between them, where such relation needs to appear in the IMPRESSION. Therefore, to better generate IMPRESSION, in addition to using important words, it is required to recognize the relations among such words in the FINDINGS and describe their corresponding relations for AIG.

In this paper, we propose to enhance AIG via a summarization model integrated with a word graph by leveraging salient words and their relations in the FINDINGS. In detail, the word graph is constructed by identifying the important words in the FINDINGS and building connections among them via different typed relations. To exploit the word graph, a Word Graph guided Summarization model (WGSUM) is designed to perform AIG, where the information from the word graph is integrated into the backbone decoder (e.g., LSTM (See et al., 2017) or Transformer (Liu and Lapata, 2019)) from two aspects, including enriching the decoder input as extra knowledge, as well as guiding the decoder to update its hidden states. Experimental results illustrate that WGSUM outperforms all baselines on two benchmark datasets, where the state-of-the-art performance is observed on all datasets in com-

parison with previous studies. Further analyses also investigate how different types of edges in the graph affects the performance of our proposed model.

2 The Proposed Method

We follow the standard sequence-to-sequence paradigm for AIG. In doing so, we regard the FINDINGS as the source sequence $\mathbf{X} = \{x_1, \dots, x_i, \dots, x_N\}$, where N is the number of tokens in the FINDINGS, and the goal of the task is to generate a target sequence (i.e., the IMPRESSION) $\mathbf{Y} = \{y_1, \dots, y_i, \dots, y_L\}$, where L is the number of tokens in the IMPRESSION. An overview of our method is shown in Figure 2 where the details are illustrated in following subsections.

2.1 The Overall Structure

The model used in our method contains three major components, i.e., the FINDINGS encoder, the graph encoder, and the graph guided decoder with their details and the training objective described below.

FINDINGS Encoder Given a FINDINGS, denoted by \mathbf{X} with N tokens, LSTM or the standard encoder from Transformer is applied to model the sequence and its output is the hidden state \mathbf{h}^x . The process is formulated as

$$\mathbf{h}^x = f_{fe}(x_1, \dots, x_i, \dots, x_N) \quad (1)$$

where $f_{fe}(\cdot)$ refers to the FINDINGS encoder.

Graph Encoder For the node V and adjacency matrix \mathbf{A} in the graph G constructed from the FINDINGS, we utilize graph neural networks² (GNN) to encode them, because GNNs are powerful in encoding graph-like information (Chen et al., 2020; Zheng and Kordjamshidi, 2020; Tian et al., 2021a,b). In detail, two encoders are employed to extract features from G , where one is used to construct the background information and the other is used to generate the dynamic guiding information. The process is thus formalized as

$$\mathbf{z}^b = f_{gb}(V, \mathbf{A}) \quad (2)$$

$$\mathbf{z}^l = f_{gl}(V, \mathbf{A}) \quad (3)$$

where $f_{gb}(\cdot)$ and $f_{gl}(\cdot)$ refer to two graph encoders, with \mathbf{z}^b and \mathbf{z}^l the intermediate states used to generate the static background information and the dynamic guiding information, respectively.

Graph Guided Decoder In our model, \mathbf{z}^b and \mathbf{z}^l from the graph encoders are integrated into the backbone decoder (e.g., LSTM and Transformer decoder.) and perform a decoding process via

$$P_{\text{vocab}} = f_d(\mathbf{h}^x, \mathbf{z}^b, \mathbf{z}^l) \quad (4)$$

where $f_d(\cdot)$ represents the decoder.

Objective Since the FINDINGS and the IMPRESSION are highly related, a pointer generator (PG) is also introduced to our model by

$$P(y_t | \mathbf{X}, y_{<t}) = p_{\text{gen}} P_{\text{vocab}}(y_t) + (1 - p_{\text{gen}}) \sum_{i: x_i = y_t} a_i^t \quad (5)$$

where \mathbf{a}_i^t is the distribution over source tokens at step t , which is obtained by performing the attention mechanism on the source tokens; p_{gen} and $(1 - p_{\text{gen}})$ are the weights of predicting the next token from the vocabulary or the source tokens, respectively. The model is then trained to maximize the negative conditional log-likelihood $P(\mathbf{Y} | \mathbf{X}, G)$ by of \mathbf{Y} given \mathbf{X} and G :

$$\theta^* = \arg \max_{\theta} \sum_{t=1}^L \log p(y_t | y_1, \dots, y_{t-1}, \mathbf{X}, G; \theta) \quad (6)$$

where θ are the parameters of the model.

2.2 Graph Construction

In FINDINGS, the most critical content that the radiologists need to summarize is the abnormal-

²It can be any type (e.g., graph convolutional networks (GCN) and graph attention networks (GAT)).

ity which usually includes the corresponding specific observations as well as their modifying words. Therefore, in our study, we first extract 5 types of entities from FINDINGS: *anatomy*, *observation*, *anatomy modifier*, *observation modifier*, and *uncertainty*, where these entities are able to cover most of the key words that need to appear in IMPRESSION. In addition, Gharebagh et al. (2020) has shown that fine-grained words are more effective than the entire ontology. Inspiring by this idea, we regard words from these entities as nodes in our graph. To avoid confusion, the repeated words are treated as one single node in each FINDINGS even though they are not presented in the same entity. Besides modifying relation, some other relations are also important for IMPRESSION generation, such as relations between *anatomy* and *observation*. For example, in Figure 1, the relation between “*pleural*” (*anatomy*) and “*effusion*” (*observation*), which can be obtained for the detailed abnormality in the IMPRESSION. To capture these types of relations, we leverage dependency trees which have been widely used to model word-word relations in many studies (Tian et al., 2020a,b; Pourn Ben Veyseh et al., 2020; Chen et al., 2021). Thus, we define three types of edges for our word graph. Note that, each FINDINGS has its corresponding word graph:

- **Type I:** this is the type using the natural order of words in an entity. In detail, we connect words if they are adjacent in the same entity. In Figure 1, the pink dashed lines serve as the type I edge. For example, “*endotracheal tube*” is an entity, so that “*endotracheal*” is connected to “*tube*”.
- **Type II:** this is the type using the relations among entities within the same category (e.g., observation and its modifier, and anatomy and its modifier). As shown in Figure 1, given an *observation*, “*effusion*” and a *observation modifier* “*moderate*”, the relation is constructed by connecting them with a green dash line.
- **Type III:** this is the type using the relations among entities across different categories (e.g., observation and anatomy). Different from the previous two types, this type is able to provide the global relation information while the previous two types emphasize the local information. In detail, we construct a dependency tree using stanza in the Universal Dependencies (UD) format (Nivre et al., 2020). As shown in Figure 1, given “*effusion*” and its head “*left*”, they are connected with a purple dash line.

Then the nodes and edges in graph are recorded by a node list V and an adjacency matrix \mathbf{A} which are then used as the input of graph encoder.

2.3 Graph Guided Decoder

We utilize graph to generate two different kinds of information and they are working on two aspects: enriching decoder input by static background information \mathbf{h}^b and controlling decoder hidden state by dynamic guiding information \mathbf{h}^l , which are introduced in following parts.

Background Information Since the graph can be considered as a condensing of source sequence (i.e., FINDINGS) which contains the most important information, it is appropriate to serve as a static background information to enrich the decoder inputs. The first output \mathbf{z}^b of graph encoder is used to construct the background information. For the hidden state \mathbf{z}_i^b in \mathbf{z}^b of each node, we can obtain attention weights by:

$$e_i^b = \mathbf{p}_b^\top \tanh(\mathbf{W}_b \mathbf{z}_i^b + \mathbf{W}_h \mathbf{h}_f) \quad (7)$$

$$\mathbf{a}^b = \text{softmax}(\mathbf{e}^b) \quad (8)$$

where $\mathbf{W}_b, \mathbf{W}_h$ and \mathbf{p}_b are learnable parameters. For LSTM, we define \mathbf{h}_f as the final hidden state, and for Transformer, we calculate the mean of all hidden states as \mathbf{h}_f . The attention distribution \mathbf{a}^b can be viewed as a probability distribution over nodes in graph. Next, \mathbf{a}^b is used to produce a weighted sum of the nodes and then we obtain the static background information:

$$\mathbf{h}^b = \sum_i a_i^b \mathbf{z}_i^b \quad (9)$$

For clarity, we simplify Equation (7), (8) and (9) as a function $\text{AttCon}(\cdot)$. Therefore, \mathbf{h}^b can be obtained by:

$$\mathbf{h}^b = \text{AttCon}(\mathbf{z}^b, \mathbf{h}_f) \quad (10)$$

In our model, the background information \mathbf{h}^b is directly concatenated to the decoder input. For each decoder input \mathbf{y}_{t-1} at step t , it is expanded as $\mathbf{y}'_{t-1} = [\mathbf{y}_{t-1}, \mathbf{h}^b]$.

Dynamic Guiding Information Since \mathbf{h}^b remains unchanged and works as the global static knowledge during the decoding process, to make the guidance more flexibly, the other information \mathbf{z}^l from graph encoder in Equation (3) is used to generate dynamic guiding information \mathbf{h}_t^l for each decoding step t . For different backbone decoders, there is a little difference in generating \mathbf{h}_t^l . For

DATA	OPENI			MIMIC-CXR		
	TRAIN	VAL	TEST	TRAIN	VAL	TEST
REPORT #	2,400	292	576	122,014	957	1,606
AFL	37.89	37.77	37.98	55.78	56.57	70.67
AFS	5.75	5.68	5.77	6.50	6.51	7.28
AFE	18	17	18	27	28	36
AIL	10.43	11.22	10.61	16.98	17.18	21.71
AIS	2.86	2.94	2.82	3.02	3.04	3.49

Table 1: The statistics of the two benchmark datasets (random split for OPENI and official split for MIMIC-CXR) in the training, validation, and test sets, where the number of reports, the averaged sentence-based (AFS, AIS) and word-based length (AFL, AIL) of IMPRESSION and FINDINGS, and the averaged number of edges in graph (AFE) are reported.

LSTM decoder, each cell updates its information by two states and one input: cell state \mathbf{c}_{t-1} , hidden state \mathbf{s}_{t-1} and input \mathbf{y}'_{t-1} , which is formulated as:

$$[\mathbf{c}_t, \mathbf{s}'_t] = \text{LSTM}(\mathbf{c}_{t-1}, \mathbf{s}_{t-1}, \mathbf{y}'_{t-1}) \quad (11)$$

where \mathbf{c}_t usually contains rich contextual information and it is appropriate to compute guiding information \mathbf{h}_t^l .³

$$\mathbf{h}_t^l = \text{AttCon}(\mathbf{z}^l, \mathbf{c}_t) \quad (12)$$

For Transformer, the general decoder only has one hidden state \mathbf{s}'_t which is the output of the last layer. In this part, we regard the output of the penultimate layer as another hidden state \mathbf{c}_t which is then used to generate dynamic information \mathbf{h}_t^l by:

$$\mathbf{h}_t^l = \text{softmax}(\mathbf{c}_t(\mathbf{z}^l)^\top) \mathbf{z}^l \quad (13)$$

After obtaining the dynamic guidance \mathbf{h}_t^l from the LSTM decoder or Transformer decoder, it is then utilized to update decoder hidden state \mathbf{s}'_t by:

$$\mathbf{s}_t = \mathbf{s}'_t + f_g(\mathbf{z}^l) \cdot f_u(\mathbf{z}^l) \quad (14)$$

where $f_g(\cdot)$ and $f_u(\cdot)$ are fully connected layers.

Vocabulary Distribution To incorporate the FINDINGS information for final prediction, we calculate the attention context vector \mathbf{g}_t^c by the same way as Equation (7), (8) and (9) using sequence encoder hidden state \mathbf{h}^x as well as the updated \mathbf{s}_t :

$$\mathbf{g}_t^c = \text{AttCon}(\mathbf{h}^x, \mathbf{s}_t) \quad (15)$$

Then both \mathbf{g}_t^c and decoder hidden state \mathbf{s}_t are used to calculate vocabulary distribution at step t :

$$P_{vocab} = \text{softmax}(\mathbf{Q}' \tanh(\mathbf{Q} [\mathbf{s}_t; \mathbf{g}_t^c])) \quad (16)$$

where \mathbf{Q}' and \mathbf{Q} are learnable weights.

³We also try to use hidden state \mathbf{s}'_t from LSTM to calculate \mathbf{h}_t^l , but the performance is not better than that of \mathbf{c}_t .

DATA	MODEL	ROUGE			FC		
		R-1	R-2	R-L	P	R	F-1
OPENI	PG-LSTM	63.21	54.13	62.78	-	-	-
	WGSUM (LSTM+GCN)	63.69	54.88	63.33	-	-	-
	WGSUM (LSTM+GAT)	64.32	55.48	63.97	-	-	-
	PG-TRANS	59.66	49.41	59.18	-	-	-
	WGSUM (Trans+GCN)	60.95	50.67	60.85	-	-	-
	WGSUM (Trans+GAT)	61.63	50.98	61.73	-	-	-
MIMIC-CXR	PG-LSTM	46.41	32.33	44.76	54.72	45.37	49.61
	WGSUM (LSTM+GCN)	46.93	32.69	45.25	55.23	46.21	50.32
	WGSUM (LSTM+GAT)	47.48	33.03	45.43	55.82	47.13	51.11
	PG-TRANS	47.16	32.31	45.47	56.18	49.08	52.39
	WGSUM (Trans+GCN)	47.93	32.63	46.23	56.37	50.84	53.46
	WGSUM (Trans+GAT)	48.37	33.34	46.68	56.83	51.22	53.88

Table 2: The performance of the baselines and our proposed methods with different GNNs on OPENI and MIMIC-CXR. R-1, R-2 and R-L denote ROUGE-1, ROUGE-2 and ROUGE-L, respectively.

3 Experiment Settings

3.1 Dataset

We conduct our experiments on the following two datasets of radiology reports: **OPENI** (Demner-Fushman et al., 2016) and **MIMIC-CXR** (Johnson et al., 2019). The former is collected by Indiana University with 3,268 reports after pre-processing. The latter contains a larger amount of data, where we obtain 124,577 after pre-processing. In our experiments, the original IMPRESSIONS written by the radiologists are considered as the ground truth. For both datasets, we follow Zhang et al. (2018) and Gharebagh et al. (2020) to exclude the following types of reports: (a) incomplete reports; (b) reports whose FINDINGS have less than 10 words; and (c) reports with IMPRESSION words less than 2. For the dataset splits, OPENI is partitioned into train/validation/test set by 2400:292:576 randomly in our experiments. For MIMIC-CXR, we apply two splits, one is the official split as the dataset published by Johnson et al. (2019), and the other is a random split with a ratio of 8:1:1, which is the same as Gharebagh et al. (2020). The statistics of the datasets are shown in Table 1.

3.2 Baseline and Evaluation Metrics

To compare the performance of our proposed models, we use the following models as our baselines:

- **PG-LSTM** (See et al., 2017): This is Pointer Generator Network (PGN) with copy mechanism where both encoder and decoder are vanilla LSTMs without graph information.
- **PG-TRANS** (Liu and Lapata, 2019): This is also a PGN with both the encoder and decoder re-

placed with the transformer.

Besides, we also compare our model with those in previous studies, including extractive summarization models, e.g., **LEXRANK** (Erkan and Radev, 2004), **TRANSFORMEREXT** (Liu and Lapata, 2019), as well as abstractive summarization models, e.g., **CAVC** (Song et al., 2020), **CGU** (Lin et al., 2018) and **ONTOLOGYABS** (Gharebagh et al., 2020). In our experiments, we use ROUGE metrics (Lin, 2004) to evaluate the generated IMPRESSIONS. We only report F₁ scores of ROUGE-1 (R-1), ROUGE-2 (R-2) and ROUGE-L (R-L), where R-1, R-2 are unigram and bigram overlap measuring the informativeness and R-L is the longest common sub-sequence overlap aiming to assess fluency. In addition, to evaluate the factual consistency (FC), CheXbert (Smit et al., 2020)⁴ is utilized to detect 14 observations related to diseases in reference impressions and generated impressions. Then precision, recall and F1 score are used to evaluate the performance.

3.3 Implementation Details

We employ stanza (Zhang et al., 2020)⁵, a python-based natural language processing library, to recognize named entities and get the syntactic analysis. Then we use the extracted entities and dependency tree to construct graph for each FINDINGS. We im-

⁴FC only apply to MIMIC-CXR since the CheXbert is designed for MIMIC-CXR and we obtain Chexbert from <https://github.com/stanfordmlgroup/CheXbert>

⁵Stanza provides packages to process the clinical text: <https://stanfordnlp.github.io/stanza/>.

MODEL	DATA	OPENI			MIMIC-CXR					
		RANDOM SPLIT			OFFICIAL SPLIT			RANDOM SPLIT		
		R-1	R-2	R-L	R-1	R-2	R-L	R-1	R-2	R-L
LEXRANK		14.63	4.42	14.06	18.11	7.47	16.87	-	-	-
TRANSFORMEREXT		15.58	5.28	14.42	31.00	16.55	27.49	-	-	-
CAVC		53.18	39.59	52.86	43.97	29.36	42.50	-	-	-
CGU		61.60	53.00	61.58	46.50	32.61	44.98	-	-	-
ONTOLOGYABS [†]		-	-	-	-	-	-	53.57	40.78	51.81
WGSUM (LSTM+GAT)		64.32	55.48	63.97	47.48	33.03	45.43	54.97	43.64	53.81
WGSUM (TRANS+GAT)		61.63	50.98	61.73	48.37	33.34	46.68	56.38	44.75	55.32

Table 3: Comparisons of our proposed models with previous study on the test sets of IU X-RAY and MIMIC-CXR with respect to ROUGE metric. † refers to that the results is directly cited from the original paper.

plement our model based on Zhang et al. (2018)⁶ and Liu and Lapata (2019)⁷. Since the quality of text representation plays an important role in model performance (Mikolov et al., 2013; Song et al., 2017, 2018; Peters et al., 2018; Song and Shi, 2018; Devlin et al., 2019; Joshi et al., 2020; Song et al., 2021), we try two powerful FINDINGS encoders, namely, LSTM and Transformer, which have achieved state-of-the-art results in many natural language processing tasks. For WGSUM (LSTM+GAT), we employ 2-layer GAT with hidden size of 200 as our graph encoder, 2-layer BiLSTM encoder for findings sequence with hidden size of 100 for each direction and 1-layer LSTM for decoder with hidden size of 200. The dropout is set to 0.5 for embedding layer. We use Adam optimizer (Kingma and Ba, 2014) with the learning rate of 0.001. For WGSUM(TRANS+GAT), the graph encoder is a 2-layer GAT with hidden size 512 and the FINDINGS encoder is a 6-layer Transformer with 512 hidden size and 2,048 feed-forward filter size. The decoder is also a 6-layer Transformer with hidden size 512.

4 Results and Analyses

4.1 Effect of word graph

To illustrate the validity of the word graph, we conduct experiments with the aforementioned baselines on the two benchmark datasets. Besides, we also try two different types of GNNs: GCN (Kipf and Welling, 2016) and GAT (Veličković et al., 2017) respectively. The results are shown in Table 2 and there are several observations drawn from different aspects. First, we can observe that for

⁶<https://github.com/yuhaozhang/summarize-radiology-findings>

⁷<https://github.com/nlpyang/PreSumm>

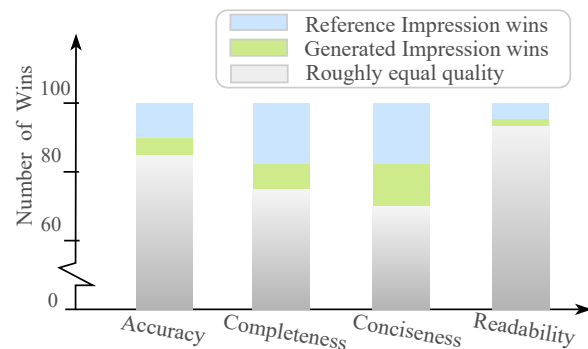


Figure 3: Experts evaluation results of 100 sampled test examples. For each of the metrics, more than 80% of the generated IMPRESSIONS are at least as good as the reference IMPRESSIONS.

our word graph, the encoder GAT is more effective than GCN where GAT can bring more significant improvements on the two datasets. The main reason might be that GAT is more powerful in updating node representation via self-attention. Second, integrating the word graph into the two different PGNs gains better performance on both the datasets, which confirms the usefulness of the word graph. Third, for OPENI, the LSTM-based models outperform much more than the Transformer-based models, while on MIMIC-CXR, the Transformer-based models are more effective. The main reason could be that the LSTM is able to obtain prominent performance in the small dataset and the Transformer is more powerful under a large amount of data. Fourth, on the FC metrics on MIMIC-CXR, our proposed methods also outperform the BASE model, indicating that the generated IMPRESSIONS from our methods are more accurate and reasonable, which is because the word graph can provide both key word and relation information for the generation process so that the decoder tends to produce words with correct relations.

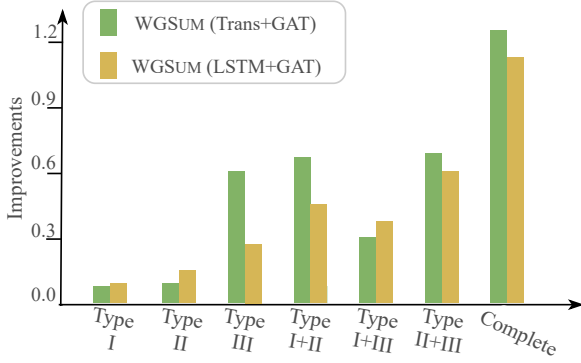


Figure 4: The improvements (R-1) of WGSUM with different graph edges on MIMIC-CXR dataset.

4.2 Comparison with Previous Studies

In this subsection, we compare our models with existing studies on the two datasets and report all results (i.e., ROUGE scores) in Table 3. We can get several observations. First, the abstractive models are apparently more effective than the extractive models in AIG, owing to the characteristics of FINDINGS and IMPRESSION in radiology reports. Second, our models with the word graph show the effectiveness of both key words information and their relations in this task when being compared to the previous models that only leverage medical term information, e.g., ONTOLOGYABS only uses ontology information in database RadLex⁸. Third, our methods achieve the best performance among all previous models, which demonstrates that using background knowledge and dynamic guidance information to control the decoding process is an appropriate design to improve the quality of the generated IMPRESSIONS.

4.3 Expert Evaluation

Since the limitation of ROUGE metrics, we further conduct an expert evaluation for a better understanding of the generated IMPRESSIONS. We randomly select 100 generated IMPRESSIONS along with their corresponding reference IMPRESSIONS and FINDINGS from MIMIC-CXR. To avoid potential bias, we randomly order the predicted and reference IMPRESSIONS. We extend Gharebagh et al. (2020) metrics to four metrics: Accuracy, Completeness, Conciseness and Readability. Three medical experts are employed to score each IMPRESSION on these metrics. Figure 3 presents the results of human evaluation. We can observe that although reference IMPRESSIONS written by the

⁸<http://www.radlex.org/Files/radlex3.10.xlsx>

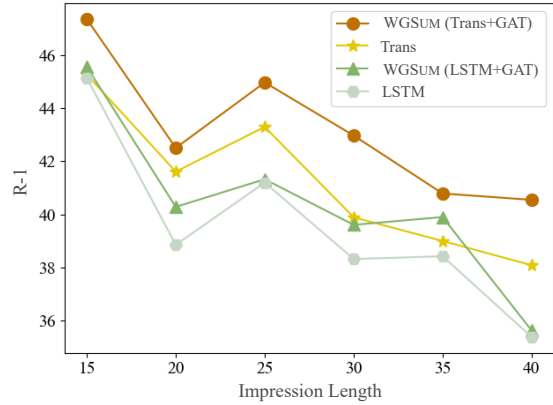


Figure 5: The performance (R-1) of generated IMPRESSIONS from WGSUM (TRANS+GAT), PG-TRANS, WGSUM (LSTM+GAT) and PG-LSTM on the MIMIC-CXR test set. Each group contains IMPRESSIONS of different length intervals.

radiologists are still better, there are still over 80% of generated IMPRESSIONS have roughly equal or better quality. About 85%, 75%, 70%, and 94% of generated IMPRESSIONS are equal to human written IMPRESSIONS on the four metrics (Accuracy, Completeness, Conciseness, and Readability). In addition, there are even 5%, 7%, 12%, 2% of generated IMPRESSIONS surpassed the reference IMPRESSIONS on these metrics.

4.4 Analyses

We conduct further analyses on Graph Edge, IMPRESSION Length, and Case Study.

Graph Edge As we introduced before, our graph contains three types of edges, i.e., entity interval edge (Type I), entity modifier edge (Type II), and edge from dependency tree (Type III). To show the effect of different edges, we conduct experiments for WGSUM (LSTM+GAT) and WGSUM (Trans+GAT) with different edges on MIMIC-CXR. The improvements from these different edge combinations are shown in Figure 4. First, we can observe that models incorporating word graph outperform the baselines no matter with what type of edge, indicating the effectiveness of our innovation in combining the entity words and their relations into the word graph. Second, regardless of WGSUM (LSTM+GAT) or WGSUM (Trans+GAT), Type III edge can bring the most significant improvements, while Type I brings little improvements. The main reason might be that the dependency tree contains more comprehensive and accurate relations for the entity words from FINDINGS. Third, WGSUM (Trans+GAT) usually obtains bet-

<p>Reference</p> <p><i>slight interval worsening of atelectasis at the left lung base. stable moderate bilateral pleural effusions , left greater than right.</i></p>	<p>Reference</p> <p><i>study yesterday , retrocardiac opacity and small opacity in right lower medial and infrahilar region has improved suggesting it was atelectasis or aspiration . there is no pleural abnormality .</i></p>	<p>Reference</p> <p><i>right internal jugular central venous catheter tip in the svc . no interval change in mild pulmonary edema with continued left basilar consolidation possibly reflecting atelectasis or infection , with small bilateral pleural effusions.</i></p>
<p>Trans</p> <p><i>1. et tube terminates 3.9 cm above the carina . 2. slight interval increase in the bibasilar atelectasis.</i></p>	<p>Trans</p> <p><i>retrocardiac opacity , likely atelectasis or aspiration .</i></p>	<p>Trans</p> <p><i>right internal jugular central venous catheter tip in the svc . no pneumothorax.</i></p>
<p>WGSUM (Trans+GAT)</p> <p><i>1. et tube terminates 3.9 cm above the carina . 2. slight bilateral pleural effusions, left greater than right , with interval worsening of bibasilar atelectasis.</i></p>	<p>WGSUM (Trans+GAT)</p> <p><i>since yesterday , retrocardiac opacity and small opacity in the right media and right lower medial lung are much better , likely atelectasis or aspiration.</i></p>	<p>WGSUM (Trans+GAT)</p> <p><i>right internal jugular central venous catheter tip in the svc. no pneumothorax. mild pulmonary edema and small bilateral pleural effusions .</i></p>

Figure 6: Examples of the generated IMPRESSIONS from two models (i.e., the PG-TRANS and the WGSUM (TRANS+GAT), respectively), as well as the reference IMPRESSIONS.

ter improvements than WGSUM (LSTM+GAT) in most cases.

IMPRESSION Length Another factor that could affect the model performance is the number of tokens in the IMPRESSION. To test the effect of IMPRESSION length, we categorize all generated IMPRESSIONS in the MIMIC-CXR test set into 6 groups (within [15,40] with the interval of 5) and compare the R-1 score for each group. The results are shown in Figure 5. There are several observations. First, when IMPRESSION gradually increases in the length, the performance of all models shows a downward trend, which indicates that long IMPRESSION generation is difficult for all models. Second, TRANSFORMER-based methods are more effective than LSTM based models, especially for long IMPRESSION. The main reason might be that the Transformer is more powerful in dealing with longer sequences via its self-attention mechanism. Third, both WGSUM (TRANS+GAT) and WGSUM (LSTM+GAT) show their superiority when being compared to their baselines and obtain better results in almost all of the groups.

Case Study To further analyze the effect of our proposed model, we perform qualitative analysis on some cases with their reference and generated IMPRESSIONS from different models. Figure 6 shows three cases from MIMIC-CXR where different colors on text refer to varied key information. It is found that in the first case, when referring to the corresponding FINDINGS our model generates more complete IMPRESSION than the reference, e.g., “*et tube terminates 3.9 cm above the carina*” is a helpful text piece but does not appear in reference IMPRESSION. In addition, compared to the reference IMPRESSIONS written by radiologists,

our method covers almost all of the key information in the generated IMPRESSIONS. For example, the key information “*moderate bilateral pleural effusions*,” “*mild pulmonary edema*” and “*small opacity in the right media*” in the three examples are not generated in PG-TRANS model, but they are necessary for describing the clinical condition.

5 Related Work

Our work focuses on summarizing the FINDINGS of radiology reports to generate the IMPRESSION, which is essentially an abstractive summarization task. For abstractive summarization, there exists a serious problem known as hallucination (Maynez et al., 2020), in which the generated summary contains fictional content. This problem also exists in AIG and would lead to misdiagnosis for the patient. To tackle this problem, in the general domain, many attempts have been made in terms of guiding information to control the generation process and output the high-quality summary (Li et al., 2018; Hsu et al., 2018; Pilault et al., 2020; Huang et al., 2020; Haonan et al., 2020). For the IMPRESSION generation task in the medical domain, there also exist several solutions. Zhang et al. (2018) encodes a section of the radiology report as the background information to guide the decoding process. MacAvaney et al. (2019) employs the entire ontological terms extracted from FINDINGS as the medical terms, and then enhances the summarizer by selecting the salient information. Gharebagh et al. (2020) further splits ontological terms into words and then incorporates these words into summarization by a separate encoder. Compared to these studies, our model offers an alternative solution to robustly enhancing guidance with a word graph for summarizing the FINDINGS of radiology reports

without requiring external resources. To our best knowledge, this is the first work employing word relation information for AIG.

6 Conclusion

In this paper, we propose a novel method for AIG, where a word graph is constructed from the FINDINGS by identifying the salient words and their relations and a graph-based model WGSUM is designed to generate IMPRESSIONS with the help of the word graph. In doing so, the information from the word graph guides the decoding process with the help of background information and dynamic guiding information. Experimental results on two benchmark datasets show the validity of our proposed method, which obtains the state-of-the-art performance on both datasets. Further analyses on the effect of edge types demonstrate that our model can generate IMPRESSION with accurate medical items.

Acknowledgements

This work is supported by Chinese Key-Area Research and Development Program of Guangdong Province (2020B0101350001) and also partially supported by NSFC under the project “The Essential Algorithms and Technologies for Standardized Analytics of Clinical Texts” (12026610). We thank three reviewers from Sun Yat-sen University: Mei Feng, Xiang Chuqi, and Liang Feiyan.

References

- Guimin Chen, Yuanhe Tian, and Yan Song. 2020. Joint Aspect Extraction and Sentiment Analysis with Directional Graph Convolutional Networks. In *Proceedings of the 28th International Conference on Computational Linguistics*, pages 272–279.
- Guimin Chen, Yuanhe Tian, Yan Song, and Xiang Wan. 2021. Relation Extraction with Type-aware Map Memories of Word Dependencies. In *Findings of the Association for Computational Linguistics: ACL-IJCNLP 2021*.
- Dina Demner-Fushman, Marc D Kohli, Marc B Rosenman, Sonya E Shooshan, Laritza Rodriguez, Sameer Antani, George R Thoma, and Clement J McDonald. 2016. Preparing a Collection of Radiology Examinations for Distribution and Retrieval. *Journal of the American Medical Informatics Association*, 23(2):304–310.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4171–4186.
- Günes Erkan and Dragomir R Radev. 2004. Lexrank: Graph-based Lexical Centrality as Saliency in Text Summarization. *Journal of artificial intelligence research*, 22:457–479.
- Sajad Sotudeh Gharebagh, Nazli Goharian, and Ross Filice. 2020. Attend to Medical Ontologies: Content Selection for Clinical Abstractive Summarization. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 1899–1905.
- Wang Haonan, Gao Yang, Bai Yu, Mirella Lapata, and Huang Heyan. 2020. Exploring Explainable Selection to Control Abstractive Generation. *arXiv*, pages arXiv–2004.
- Saeed Hassanpour and Curtis P Langlotz. 2016. Information Extraction from Multi-institutional Radiology Reports. *Artificial intelligence in medicine*, 66:29–39.
- Wan-Ting Hsu, Chieh-Kai Lin, Ming-Ying Lee, Kerui Min, Jing Tang, and Min Sun. 2018. 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)*, pages 132–141.
- Luyang Huang, Lingfei Wu, and Lu Wang. 2020. Knowledge Graph-Augmented Abstractive Summarization with Semantic-Driven Cloze Reward. *arXiv preprint arXiv:2005.01159*.
- Alistair EW Johnson, Tom J Pollard, Nathaniel R Greenbaum, Matthew P Lungren, Chih-ying Deng, Yifan Peng, Zhiyong Lu, Roger G Mark, Seth J Berkowitz, and Steven Horng. 2019. MIMIC-CXR-JPG, a Large Publicly Available Database of Labeled Chest Radiographs. *arXiv preprint arXiv:1901.07042*.
- Mandar Joshi, Danqi Chen, Yinhan Liu, Daniel S. Weld, Luke Zettlemoyer, and Omer Levy. 2020. SpanBERT: Improving pre-training by representing and predicting spans. *Transactions of the Association for Computational Linguistics*, 8:64–77.
- Diederik P Kingma and Jimmy Ba. 2014. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*.
- Thomas N Kipf and Max Welling. 2016. Semi-supervised Classification with Graph Convolutional Networks. *arXiv preprint arXiv:1609.02907*.

- Chenliang Li, Weiran Xu, Si Li, and Sheng Gao. 2018. Guiding Generation for Abstractive Text Summarization Based on Key Information Guide Network. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 2 (Short Papers)*, pages 55–60.
- Chin-Yew Lin. 2004. Rouge: A Rackage for Automatic Evaluation of Summaries. In *Text summarization branches out*, pages 74–81.
- Junyang Lin, Xu Sun, Shuming Ma, and Qi Su. 2018. Global Encoding for Abstractive Summarization. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 163–169.
- 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.
- Sean MacAvaney, Sajad Sotudeh, Arman Cohan, Nazli Goharian, Ish Talati, and Ross W Filice. 2019. Ontology-Aware Clinical Abstractive Summarization. In *Proceedings of the 42nd International ACM SIGIR Conference on Research and Development in Information Retrieval*, pages 1013–1016. ACM.
- Joshua Maynez, Shashi Narayan, Bernd Bohnet, and Ryan McDonald. 2020. On Faithfulness and Factuality in Abstractive Summarization. *arXiv*, pages arXiv–2005.
- Tomas Mikolov, Kai Chen, Greg Corrado, and Jeffrey Dean. 2013. Efficient Estimation of Word Representations in Vector Space. *arXiv preprint arXiv:1301.3781*.
- Joakim Nivre, Marie-Catherine de Marneffe, Filip Ginter, Jan Hajič, Christopher D Manning, Sampo Pyysalo, Sebastian Schuster, Francis Tyers, and Daniel Zeman. 2020. Universal Dependencies v2: An Evergrowing Multilingual Treebank Collection. *arXiv preprint arXiv:2004.10643*.
- Matthew Peters, Mark Neumann, Mohit Iyyer, Matt Gardner, Christopher Clark, Kenton Lee, and Luke Zettlemoyer. 2018. Deep contextualized word representations. 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 2227–2237, New Orleans, Louisiana.
- Jonathan Pilault, Raymond Li, Sandeep Subramanian, and Christopher Pal. 2020. On Extractive and Abstractive Neural Document Summarization with Transformer Language Models. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 9308–9319.
- Amir Pouran Ben Veysseh, Nasim Nouri, Franck Dernoncourt, Quan Hung Tran, Dejing Dou, and Thien Huu Nguyen. 2020. Improving Aspect-based Sentiment Analysis with Gated Graph Convolutional Networks and Syntax-based Regulation. In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 4543–4548, Online.
- Abigail See, Peter J Liu, and Christopher D Manning. 2017. 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)*, pages 1073–1083.
- Akshay Smit, Saahil Jain, Pranav Rajpurkar, Anuj Pareek, Andrew Y Ng, and Matthew Lungren. 2020. Chexbert: Combining Automatic Labelers and Expert Annotations for Accurate Radiology Report Labeling Using BERT. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 1500–1519.
- Kaiqiang Song, Bingqing Wang, Zhe Feng, Ren Liu, and Fei Liu. 2020. Controlling the Amount of Verbatim Copying in Abstractive Summarization. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, pages 8902–8909.
- Yan Song, Chia-Jung Lee, and Fei Xia. 2017. Learning Word Representations with Regularization from Prior Knowledge. In *Proceedings of the 21st Conference on Computational Natural Language Learning (CoNLL 2017)*, pages 143–152.
- Yan Song and Shuming Shi. 2018. Complementary Learning of Word Embeddings. In *Proceedings of the Twenty-Seventh International Joint Conference on Artificial Intelligence, IJCAI-18*, pages 4368–4374.
- Yan Song, Shuming Shi, Jing Li, and Haisong Zhang. 2018. Directional Skip-Gram: Explicitly Distinguishing Left and Right Context for Word Embeddings. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 2 (Short Papers)*, pages 175–180.
- Yan Song, Tong Zhang, Yonggang Wang, and Kai-Fu Lee. 2021. ZEN 2.0: Continue Training and Adaption for N-gram Enhanced Text Encoders. *arXiv preprint arXiv:2105.01279*.
- Yuanhe Tian, Guimin Chen, and Yan Song. 2021a. Aspect-based Sentiment Analysis with Type-aware Graph Convolutional Networks and Layer Ensemble. In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 2910–2922, Online.
- Yuanhe Tian, Guimin Chen, Yan Song, and Xiang Wan. 2021b. Dependency-driven Relation Extraction with Attentive Graph Convolutional Networks.

In *Proceedings of the Joint Conference of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing*.

Yuanhe Tian, Wang Shen, Yan Song, Fei Xia, Min He, and Kenli Li. 2020a. Improving Biomedical Named Entity Recognition with Syntactic Information. *BMC Bioinformatics*, 21:1471–2105.

Yuanhe Tian, Yan Song, and Fei Xia. 2020b. Supertagging Combinatory Categorical Grammar with Attentive Graph Convolutional Networks. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 6037–6044.

Petar Veličković, Guillem Cucurull, Arantxa Casanova, Adriana Romero, Pietro Lio, and Yoshua Bengio. 2017. Graph Attention Networks. *arXiv preprint arXiv:1710.10903*.

Yuhao Zhang, Daisy Yi Ding, Tianpei Qian, Christopher D Manning, and Curtis P Langlotz. 2018. Learning to Summarize Radiology Findings. *EMNLP 2018*, page 204.

Yuhao Zhang, Yuhui Zhang, Peng Qi, Christopher D Manning, and Curtis P Langlotz. 2020. Biomedical and Clinical English Model Packages in the Stanza Python NLP Library. *arXiv preprint arXiv:2007.14640*.

Chen Zheng and Parisa Kordjamshidi. 2020. SRLGRN: Semantic role labeling graph reasoning network. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 8881–8891, Online.