

MEPNet: Medical Entity-Balanced Prompting Network for Brain CT Report Generation

Xiaodan Zhang¹, Yanzhao Shi¹, Junzhong Ji^{1*}, Chengxin Zheng¹, Liangqiong Qu^{2*}

¹College of Computer Science, Beijing University of Technology, Beijing, China

²School of Computing and Data Science, The University of Hong Kong, Hong Kong, China

zhangxiaodan@bjut.edu.cn, yanzhaoshi0927@outlook.com,

jjz01@bjut.edu.cn, chaunxey_zeta@emails.bjut.edu.cn, liangqu@hku.hk

Abstract

The automatic generation of brain CT reports has gained widespread attention, given its potential to assist radiologists in diagnosing cranial diseases. However, brain CT scans involve extensive medical entities, such as diverse anatomy regions and lesions, exhibiting highly inconsistent spatial patterns in 3D volumetric space. This leads to biased learning of medical entities in existing methods, resulting in repetitiveness and inaccuracy in generated reports. To this end, we propose a Medical Entity-balanced Prompting Network (MEPNet), which harnesses the large language model (LLM) to fairly interpret various entities for accurate brain CT report generation. By introducing the *visual embedding* and the *learning status* of medical entities as enriched clues, our method prompts the LLM to balance the learning of diverse entities, thereby enhancing reports with comprehensive findings. First, to extract visual embedding of entities, we propose Knowledge-driven Joint Attention to explore and distill entity patterns using both explicit and implicit medical knowledge. Then, a Learning Status Scorer is designed to evaluate the learning of entity visual embeddings, resulting in unique learning status for individual entities. Finally, these entity visual embeddings and status are elaborately integrated into multi-modal prompts, to guide the text generation of LLM. This process allows LLM to self-adapt the learning process for biased-fitted entities, thereby covering detailed findings in generated reports. We conduct experiments on two brain CT report generation benchmarks, showing the effectiveness in clinical accuracy and text coherence.

Code — <https://github.com/YanzhaoShi/MEPNet>.

Introduction

Radiology reports provide interpretations of detailed pathological findings from medical images, which are crucial for patient treatments. Despite the significance, for radiologists, hand-crafted writing practices are time-consuming and may cause error outcomes due to subjective factors (e.g. diversion and exhaustion) (Brady et al. 2012). The advent of automated medical report models can produce accurate reports and offer potential benefits for medical examinations, such as reducing the workload shouldered by physicians and saving the limited clinical resources in densely populated countries.

* Corresponding Authors

Copyright © 2025, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

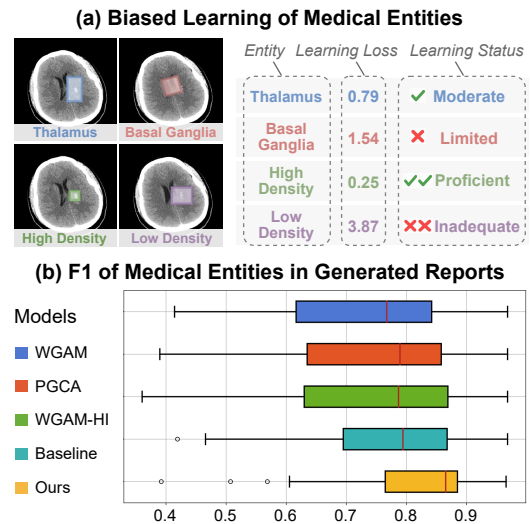


Figure 1: **(a)** shows an example of biased entity learning, where different medical entities in brain CT scans exhibit distinct learning losses, indicating biased learning statuses. **(b)** compares F1 scores of different MRG models for covering various medical entities in generated reports. Our method shows the narrowest box with a higher median F1 score, indicating our effectiveness in balanced and accurate learning.

In recent years, medical report generation (MRG) (Jing, Xie, and Xing 2018; Yang et al. 2021, 2023; Shi et al. 2023; Li et al. 2023) has emerged as a core topic within the medical multi-modal learning. Different from natural image captioning (Xu et al. 2015; Ji et al. 2020), MRG models focus on reporting tiny pathological conditions from sparse and similar visual information. This motivates MRG models to learn specific patterns associated with *medical entities*, encompassing entities of *anatomy regions* such as “basal ganglia” and entities of *lesions* such as “high density” (refer to Figure 1(a)). To achieve this, task-specific attention mechanisms (Jing, Xie, and Xing 2018; Wang et al. 2018) are early designed to ground medical semantics with specific visual areas. Memory mechanism (Chen et al. 2021; Qin and Song 2022) is then proposed to store cross-modal patterns of diseases for generating accurate content. Moreover, knowledge graph (Liu

et al. 2021; Li et al. 2023; Shi et al. 2023) is another effective approach for MRG models to grasp prior relations of medical entities. Recently, as large language models (LLMs) showcase impressive capabilities in language processing, a growing number of studies have started to employ LLMs as decoders for MRG. Through carefully designing medical prompts (Bu et al. 2024; Chen et al. 2024) and fine-tuning strategies (He et al. 2024), these methods have been proven to be more effective than traditional MRG models.

Despite the burgeoning popularity of LLM-based methods, the focus of most existing models remains on 2D imaging, e.g. chest X-rays. In contrast, brain CT data encompasses more complex pathological conditions in 3D volumetric space, with a broader range of medical entities (Shi et al. 2023). These medical entities display diverse visual distributions, grounded by unique shapes of cerebral organs or contours of intracranial lesions (Bhadoria and Dewal 2014). This visual diversity is vital for radiologists to identify and diagnose. However, empowering LLMs to effectively capture these diverse visual clues is challenging (Chen et al. 2024; Blankemeier et al. 2024), which results in a huge learning bias across different entities (see Figure 1(a)). This bias disperses the language model’s focus towards entities that should not be overly emphasized or ignored, leading to the generation of repetitive and inaccurate diagnostic sentences. Although logit-adjusted classifier (Jin et al. 2024) is recently proposed to mitigate this bias by prompting the language model with well-classified textual labels of entities, it reduces the LLM’s sensitivity to interpret visual patterns since it relies too heavily and too early on classified texts. Thus, how to prompt LLMs to fairly grasp visual patterns across various medical entities remains a critical concern in MRG tasks.

In this paper, we propose a Medical Entity-balanced Prompting Network (MEPNet) to enhance the capability of LLM in balancing the learning of diverse medical entities for high-quality brain CT report generation. The core idea is to introduce the *visual embedding* and *learning status* for individual medical entities, which are integrated as enriched clues to prompt LLM to fairly understand entities’ visual patterns. Specifically, the learning status is employed to evaluate the quality of entity visual embedding. To derive visual embedding for entities, we propose a Knowledge-driven Joint Attention. It captures and distills entity patterns from CT scans via explicit medical knowledge (summarized from a knowledge graph developed under expert guidance) and implicit medical knowledge (explored by the model itself). Next, a Learning Status Scorer is introduced to evaluate the learning statuses of entities, which are then converted to status words, e.g. *limited*, *moderate*, and *proficient*, to represent varying degrees of learning. Finally, we integrate these entity visual embeddings and related status as entity-balanced prompts, which are elaborately combined with scan visual embeddings and task-specific instructions to form multi-modal prompts, guiding LLM in generating texts. As evidenced in Figure 1(b), our method achieves a balance in precisely generating diverse medical entities in reports.

Our main contributions can be summarized as:

1. We propose a novel method that balances the LLM in learning diverse medical entities for accurate brain CT

report generation. This is achieved by integrating enriched medical clues, which include visual embeddings and related learning status of entities, to prompt LLM fairly generate findings.

2. We design a Knowledge-driven Joint Attention to exploit high-level visual embeddings of medical entities from 3D volumetric space by using explicit and implicit medical knowledge; We further introduce a Learning Status Scorer to evaluate the quality of entity visual embeddings.
3. We comprehensively evaluate our model on the BCT-CHR and CTRG datasets. The results confirm the effectiveness of our method in generating accurate brain CT reports.

Related Works

With the development of intelligent healthcare, medical report generation (MRG), as an essential yet challenging application, has received notable attention (Jing, Xie, and Xing 2018; Li et al. 2019; Jia et al. 2020; Zhang et al. 2020; Song et al. 2022; Zhang et al. 2023a; Shen et al. 2024; Chen et al. 2024; Zheng et al. 2024). To satisfy the accuracy in generating long medical descriptions, existing efforts can be roughly divided into three types. The first type is enhancing visual features via domain knowledge for producing precise texts. Jing, Xie, and Xing (2018) propose to inject medical concepts into visual embeddings by a co-attention mechanism. Li et al. (2019) and Zhang et al. (2020) mine disease relations from graph knowledge and use them to enhance visual contexts. The second type of MRG is to conduct visual-textual alignment to improve the cross-modal correspondence. To achieve this, memory mechanism (Chen et al. 2021; Qin and Song 2022; Wang, Bhalerao, and He 2022) is designed to map multi-modal features into mutual space for generating accurate sentences grounded by relevant visual patterns. Contrastive learning (Yan et al. 2021; Shi et al. 2023; Yang et al. 2023; Li et al. 2023; Shi et al. 2024) is also widely applied to model cross-modal representations in an unsupervised manner, showing efficiency in small-scale medical datasets. The recent advance in LLMs inspires the third type of MRG, which can replace previous complex model designs with unified LLMs while improving task performance (Zhou et al. 2024). Yan et al. (2023) froze an LLM and leverage in-context learning to generate human-style chest X-ray reports. Dia-LLaMA (Chen et al. 2024) fine-tunes LLaMA2-7B (Touvron et al. 2023) by using discriminated disease labels as knowledge prompts. Bu et al. (2024) extract knowledge of pulmonary lesions from patient instances to facilitate X-ray report generation. Liu et al. (2024a) further bootstrap LLMs for MRG by instance induction and coarse-to-fine decoding. Compared with them, we focus on more complex 3D brain CT scans and offer a novel prompting paradigm to represent medical entity clues with informative embeddings and learning status. This activates the inner power of LLM to balance entity learning and reduce the hallucinations in generating diagnostic sentences.

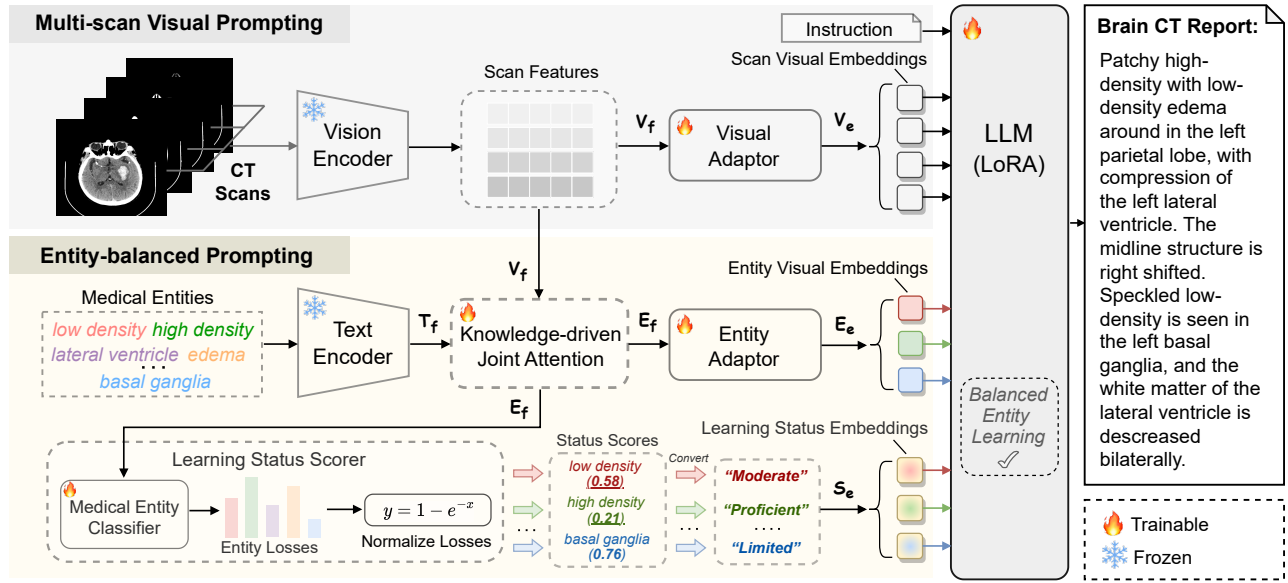


Figure 2: The framework of MEPNet, with two branches: Multi-scan Visual Prompting, which processes visual information of brain CT scans, and Entity-balanced Prompting, which mines entity visual embeddings and corresponding learning status for achieving balanced entity learning within LLM. These branches collaboratively prompt the LLM to generate diagnostic reports.

Methods

3.1 Overall Framework

An overview of MEPNet is shown in Figure 2. Given a patient sample with N brain CT scans $S = \{s_1, \dots, s_N\}$, the target is to generate a brain CT report $R = \{r_1, \dots, r_M\}$ with M textual words. We employ LLM as the decoder for report generation, with enriched prompts derived from two distinct branches: The multi-scan visual prompting branch and the entity-balanced prompting branch.

Multi-scan Visual Prompting To obtain visual features from volumetric CT scans, we follow Zhang et al. (2023b) and utilize a finetuned ResNet101 model to extract scan features $V_f = \{v_{f_1}, \dots, v_{f_N}\} \in \mathbb{R}^{N \times P \times d_s}$ for each sample, where P and d_s denote the patch number and feature channels, respectively. We then employ a visual adaptor to convert V_f into the word embedding space d_w within LLM:

$$V_e = (V_f W_{v_1}^\top + b_{v_1}) W_{v_2}^\top + b_{v_2}, \quad (1)$$

where $V_e \in \mathbb{R}^{N \times d_w}$ denotes the scan visual embedding, $W_{v_1} \in \mathbb{R}^{(P \cdot d_s) \times d_s}$ and $W_{v_2} \in \mathbb{R}^{d_s \times d_w}$ are learnable weights, b_{v_1} and b_{v_2} and are learnable biases.

Entity-Balanced Prompting Adding medical information into prompts is proven to enhance the diagnostic capabilities of LLMs (Jin et al. 2024). Unlike previous methods that construct text prompts based on classified categories, we propose to use informative visual embeddings of medical entities as prompts. This forces the LLM to grasp medical patterns from redundant visual features, rather than just taking shortcuts by rewriting the classified entity labels. We first collect a set of medical entities $E = \{e_0, e_1, \dots, e_k\}$, where K entities $e_{1:k}$ are summarized by Shi et al. (2023) and e_0 denotes a global

entity (i.e. $[global]$) to represent overall conditions. We extract textual features of $k + 1$ entities as $T_f \in \mathbb{R}^{(k+1) \times d_w}$ via LLM’s text embedding layer. Then, T_f is used to capture detailed entity visual features $E_f \in \mathbb{R}^{(k+1) \times d_h}$ from the whole scan features V_f via Knowledge-driven Joint Attention (KJA) module (see Section 3.2.). Next, we remove the global entity from E_f and map the remaining feature into LLM’s embedding space, resulting in entity visual embeddings $E_e \in \mathbb{R}^{k \times d_w}$.

Additionally, we introduce a learning status scorer to evaluate the learning status of k entities, which is denoted as $T_s = \{t_{s_1}, \dots, t_{s_k}\} \in \mathbb{R}^k$. We convert T_s into status words $S_w = \{s_{w_1}, \dots, s_{w_k}\}$, which is embedded as learning status embeddings S_e to depict the learning degree of each entity.

Model Training We integrate the scan visual embedding V_e , entity visual embedding E_e , and learning status embedding S_e into multi-modal prompt M_P , which guides a LLM-based decoder to generate reports. The template of our multi-modal prompt is shown in Figure 3, where V_e , E_e , and S_e are denoted as $\{scan\}$, $\{entity_embedding\}$, and $\{entity_status\}$, respectively. The remaining contents are instructions, where $\{entity\ i\}$ is the entity word e_i . We also add special tokens, e.g. $[Img]$ and $[/Img]$ for visual input identification, and $[MRG]$ for task identification. We finetune LLM by LoRA (Hu et al. 2021) and use cross-entropy loss for optimization:

$$\mathcal{L}_g = - \sum_{t=1}^M \log P(y_t | y_{1:t-1}, M_P; \theta) \quad (2)$$

where $P(y_t | *)$ denotes the probability conditioned on multi-modal prompt M_P and the embeddings of previous words y_1, y_2, \dots, y_{t-1} . θ denotes the trainable parameters in LLM.

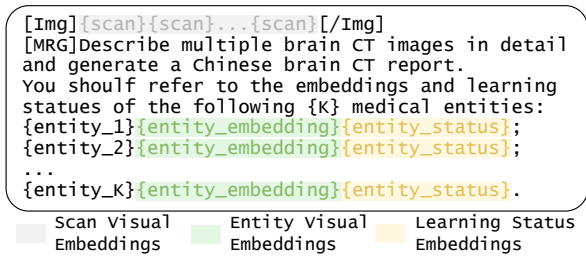


Figure 3: Template of the multi-modal prompt. Different colors are utilized to represent distinct components.

3.2 Knowledge-Driven Joint Attention

Given entity textual features T_f and scan features V_f , this module aims to capture detailed visual features for entities based on entities' medical relationships. Unlike directly injecting relations into visual features (Zhang et al. 2020) or textual keywords (Li et al. 2023), we argue that they may fail to learn multi-modal medical clues. Instead, we first integrate T_f into V_f by cross attention, and then use knowledge-masked self attention to inject relations.

Cross Attention We project entity textual features T_f into d_h channels to form query vector Q_e . The channel of scan features V_f is also projected to d_h to form key vector K_c and value vector S_c . The cross attention can be formulated as:

$$T'_f = FFN(MHA(Q_e, K_c, S_c)), \quad (3)$$

where $T'_f \in \mathbb{R}^{(k+1) \times d_h}$ primarily depicts the visual features of medical entities. FFN and MHA denote the feed-forward network and multi-head attention, respectively.

Knowledge Extraction The obtained T'_f is still limited to representing complex visual patterns since it overlooks the inherent medical relations of entities that are vital for diagnosis. To tackle this, we extract entities' relations from both explicit and implicit perspectives as medical knowledge.

Explicit entity relations depict empirical medical knowledge defined by experts. We acquire this knowledge from brain pathological graph (Shi et al. 2023). This graph divides medical entities E into anatomy regions $E_r \in E$ and lesions $E_l \in E$, with three types of fine-grained edges connecting them: 1) R_{t2t} : edges among anatomy regions; 2) R_{l2l} : edges among lesions; 3) R_{t2l} : edges between anatomy regions and lesions. R_{t2t} and R_{l2l} are fixed by medical priors, and R_{t2l} is built based on the specific tissue-lesion relations in ground truth report. Note that during model inference, we do not extract R_{t2l} from ground truth report to prevent data leakage. Instead, we substitute it with the statistical information of tissue-lesion relationships derived from the training data. Based on this graph, given one entity $e_i \in E$ as a query, we can find its relations with other entities, which can be denoted as $R_{i2E} \in \mathbb{R}^{k+1}$. The relations associated with the global entities e_0 are set to 1. We extract all entities' relations and concatenate them as medical adjacency matrix $M_E \in \mathbb{R}^{(k+1) \times (k+1)}$ (see Figure 4).

Besides, there also exists some implicit yet essential relations among medical entities that haven't been explored

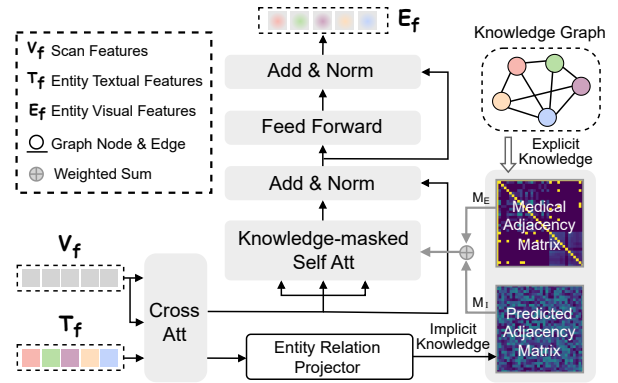


Figure 4: Details of the Knowledge-driven Joint Attention.

by radiologists. To fully consider these potential entity relations, we introduce an entity relation projector to predict them automatically, which can be formulated as:

$$R_e = W_{r_e} T'_f + b_{r_e}, \quad (4)$$

$$M_I = \text{Sigmoid}(W_{r_p} R_e R_e^T + b_{r_p}), \quad (5)$$

where $\text{Sigmoid}(\cdot)$ is the sigmoid function. $M_I \in \mathbb{R}^{(k+1) \times (k+1)}$ is the predicted adjacency matrix. W_{r_e} , W_{r_p} are learnable weight matrixes, b_{r_e} , b_{r_p} are learnable biases.

We combine two types of matrixes into the entity adjacency matrix M_{adj} , which can be formulated as:

$$M_{adj} = \alpha_E M_E + \alpha_I M_I, \quad (6)$$

where α_E and α_I are hyperparameters to balance the weights of M_E and M_I .

Knowledge-Masked Self Attention We employ M_{adj} as the attention mask and conduct knowledge-masked self attention to inject medical knowledge into T'_f :

$$e_{att} = LN(KSA(T'_f, M_{adj}) + T'_f), \quad (7)$$

$$E_f = LN(FFN(e_{att}) + e_{att}), \quad (8)$$

where KSA is masked multi-head attention with M_{adj} as mask, LN is the layer norm. $E_f = \{e_{f_0}, \dots, e_{f_k}\} \in \mathbb{R}^{(k+1) \times d_h}$ denotes the entity visual features. e_{f_i} denotes the visual feature of the i -th entity, which aggregates its neighbor entity features according to relations in M_{adj} .

3.3 Learning Status Scorer

To address the biased learning of medical entities, we propose a learning status scorer to evaluate the learning status of each entity at each training step. These learning statuses are used to prompt the LLM to balance the learning of diverse entities.

Score Generation The scorer is based on calculating discriminative loss for each medical entity. Followed by Jin et al. (2024), we design a medical entity classifier to obtain discriminative loss. The classifier can be formulated as:

$$E_h = W_{e_h} E_f + b_{e_h}, \quad (9)$$

$$E_p = \text{Sigmoid}(Cls(E_h)), \quad (10)$$

where $Cls(\cdot)$ denotes a classification head, $E_p = \{e_{p_1}, \dots, e_{p_k}\} \in \mathbb{R}^k$ denotes the predicted labels for k entities. Then, given a ground truth label $\hat{E}_p = \{\hat{e}_{p_1}, \dots, \hat{e}_{p_k}\} \in \{0, 1\}^k$, we gather labels of individual entities in a batch level to measure the loss of each entity, e.g., the loss for entity e_i can be calculated as:

$$e_{s_i} = - \sum_j^b W_{e_j} [e_{p_i}^{(j)} \log \hat{e}_{p_i}^{(j)} + (1 - e_{p_i}^{(j)}) \log(1 - \hat{e}_{p_i}^{(j)})], \quad (11)$$

where b denotes the batch size, $e_{p_i}^{(j)}$ and $\hat{e}_{p_i}^{(j)}$ denote the ground truth label and predicted label of entity e_i in the j -th sample, respectively. Then, we apply the Exponential Saturation Function $y = 1 - e^{-x}$ to normalize the loss to the range of 0 to 1. We use the normalized entity loss as the score of learning status, where a higher score indicates a worse learning of related entity. The scores for all entities are represented as $E_s = \{e_{s_1}, \dots, e_{s_k}\} \in \mathbb{R}^k$.

Meanwhile, to ensure a good evaluation of learning status, we optimize the classifier during the training:

$$\mathcal{L}_d = - \sum_i^b W_{d_i} [E_p^{(i)} \log \hat{E}_p^{(i)} + (1 - E_p^{(i)}) \log(1 - \hat{E}_p^{(i)})], \quad (12)$$

where $E_p^{(i)}$ and $\hat{E}_p^{(i)}$ represent the ground truth label and predicted label of sample i .

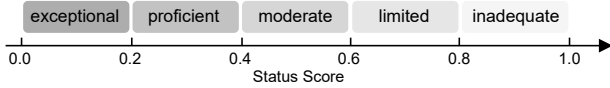


Figure 5: Matching rules of status words and status scores.

Status Word Matching To let the LLM understand the status scores for balancing entity learning, we convert the generated scores to textual words, as LLM has an inherent sensitivity to interpreting texts. As shown in Figure 5, we set 5 types of status words to match different ranges of scores. Given scores $E_s \in \mathbb{R}^k$, the converted words can be denoted as $S_w = \{s_{w_1}, \dots, s_{w_k}\} \in \mathbb{R}^k$. We then employ LLM’s text embedding layer to encode S_w into learning status embeddings $S_e \in \mathbb{R}^{k \times d_w}$. For status words that are tokenized into multiple tokens by LLM, we average their embeddings to represent the centroid semantics, which reduces computational costs while preserving the initial semantics (Li et al. 2023).

3.4 Training Objective

Our overall loss contains report generation loss and classification loss, which can be represented as:

$$\mathcal{L} = \mathcal{L}_g + \lambda \mathcal{L}_d, \quad (13)$$

where λ is the coefficient to balance the module losses.

Experiments and Results

4.1 Datasets, Metrics and Settings

BCT-CHR BCT-CHR (Yang et al. 2021) is the benchmark for brain CT report generation. It comprises 49,152 brain CT

scans paired with 2,048 Chinese reports. Each example includes 24 CT scans along with a report. We split the training, testing, and validation sets in a ratio of 7:2:1. Note that in this paper, we only give English translations of Chinese reports for better understanding.

CTRG-Brain CTRG-Brain (Tang et al. 2024) is an open-source dataset, which contains 6007 samples with cranial abnormalities and includes 160,352 brain CT scans and 6007 Chinese reports. We employ brain CT data cleaning pipeline (Zhang et al. 2023b) to filter redundant scans and ensure each sample contains 24 scans, aligning the data with medical standards. Since the official split misses a validation set, we align its data split with the BCT-CHR in 7:2:1.

Metrics We use natural language generation (NLG) and clinical evaluation (CE) metrics to evaluate the report generation. The NLG metrics contains BLEU (Papineni et al. 2002), METEOR (Lavie and Agarwal 2007), ROUGE-L (Lin 2004) and CIDER (Vedantam, Zitnick, and Parikh 2015). For CE metrics, we follow Shi et al. (2023) and collect 24 keywords, e.g. “basal ganglia region”, “cerebral sulcus” and “stagnant blood”, to calculate Precision, Recall, and F1 scores.

Settings Our baseline follows encoder-decoder architecture. We utilize ResNet101 (He et al. 2016) to process visual features, which is fine-tuned for the brain hemorrhage classification using the CQ500 dataset (Chilamkurthy et al. 2018). For the textual decoder, we utilize LLaMA3-8B (Meta 2024), configured with a 4-bit quantization setting. LoRA (Hu et al. 2021) (with a rank of 64) is also adopted for efficient fine-tuning, with 0.34% trainable parameters in LLM (27.3M). We use Chinese instructions to prompt LLM, and the template in Figure 3 is translated into English for better clarity. The feature channels $\{d_h, d_s, d_w\}$, knowledge weights $\{\alpha_E, \alpha_I\}$, and classification loss weight λ are set to $\{512, 2048, 4096\}$, $\{0.9, 0.1\}$, and 0.1, respectively. We train the model on two RTX 3090 GPUs via AdamW optimizer with a learning rate of $1e-4$ for 10 epochs, and the batch size is set to 4.

4.2 Comparison Studies

We compare our model with advanced brain CT report generation methods WGAM (Yang et al. 2021), PGCA (Shi et al. 2023), and WGAM-HI (Zhang et al. 2023b). Besides, we also reproduce some other models initially designed for image captioning (Up-Down (Anderson et al. 2018)), chest X-ray MRG (WCL (Yan et al. 2021), PromptMRG (Jin et al. 2024)), and chest CT MRG (HILT (Liu et al. 2024b)) for comprehensive comparisons.

As shown in Table 1, our model achieves superior performance on both two datasets, highlighting the effectiveness of our entity-balanced learning strategy. In contrast, Up-Down yields inferior results (especially on B4) due to the limited focus on medical information. WGAM utilizes weakly guided visual attention to learn spatial features, while WGAM-HI enhances it in a multi-scale fashion and achieves higher CE scores. Compared with traditional decoders, HILT adopts LLM as decoder, which performs even better than competitors originally designed for brain CT MRG on the CTRG-Brain dataset. Surprisingly, HILT performs worse on

Datasets	Methods	NLG Metrics				CE Metrics		
		B4	M	R-L	C	F1	Recall	Precision
CTRG-B	Up-Down(Anderson et al. 2018)†	24.4	31.6	42.5	27.3	-	-	-
	WCL(Yan et al. 2021)†	25.1	31.3	42.8	33.3	-	-	-
	WGAM(Yang et al. 2021)†	25.4	32.0	42.4	31.9	-	-	-
	PGCA(Shi et al. 2023)†	26.5	32.5	43.0	34.0	-	-	-
	WGAM-HI(Zhang et al. 2023b)†	26.1	31.4	43.8	33.2	-	-	-
	PromptMRG(Jin et al. 2024)†	26.6	31.0	47.4	50.3	-	-	-
	HILT(Liu et al. 2024b)†	33.3	33.0	50.3	80.0	-	-	-
	MEPNet (Ours)	34.4	34.6	50.4	86.6	-	-	-
BCT-CHR	Up-Down(Anderson et al. 2018)†	13.5	27.0	35.4	20.1	47.2	59.8	49.5
	WCL(Yan et al. 2021)†	14.2	26.9	35.3	18.8	42.4	50.2	51.0
	WGAM(Yang et al. 2021)†	14.5	27.8	35.1	17.3	41.6	52.6	52.8
	PGCA(Shi et al. 2023)†	15.5	28.7	36.5	19.9	49.6	61.9	52.0
	WGAM-HI(Zhang et al. 2023b)†	15.6	29.0	36.5	22.2	51.4	66.8	51.1
	PromptMRG(Jin et al. 2024)†	12.9	26.8	34.0	16.1	49.6	52.9	43.3
	HILT(Liu et al. 2024b)†	12.6	27.3	33.6	15.9	43.8	52.4	39.9
	MEPNet (Ours)	14.7	29.2	36.8	22.3	52.0	68.9	50.6

Table 1: Performance comparison of our MEPNet with state-of-the-art models on two brain CT report generation datasets. Models are categorized into those utilizing traditional decoders and those employing LLM decoders. Superior results are highlighted in bold. Note that CE metrics are solely evaluated on the BCT-CHR dataset, as the CTRG-Brain dataset lacks a unified evaluation criterion. † denotes the re-produced models.

a relatively small BCT-CHR dataset. The reason may be that without appropriate prompts to connect the visual encoder and text decoder, standard LLM struggles to efficiently learn from limited medical data. Although PromptMRG prompts the language model with classified disease labels, it overly relies on the given labels and only gains modest scores. Notably, our method excels with significant gains in CIDEr, showing the advances in handling the repetitiveness in generated reports. The SoTA score on Recall also indicates the comprehensive coverage of medical keywords. In Figure 1(b), we are encouraged to find that our method shows the largest mean and smallest variance for generating diverse entities, indicating the achievement of balanced medical entity learning. Despite the B4 rises on CTRG-Brain, the score on BCT-CHR remains lower. This may be due to the more flexible writing styles in BCT-CHR reports, as we do not regulate grammar structure or filler words in LLM, leading to sentence rephrasing and reducing multi-gram word matching. Additionally, our model has lower precision, likely because our status embeddings prompt the LLM to generate a broader range of entities, lowering precision but highly boosting recall and F1 scores, which are more beneficial for diagnosis assistance.

4.3 Ablation Studies

We conduct ablation studies on components of entity-balanced prompting (i.e. entity visual embedding E_e and status embedding S_e) and types of medical knowledge within knowledge-driven joint attention (i.e. M_E and M_I).

As shown in Table 2, our baseline performs worse than other settings, which indicates that it is difficult for LLM to directly understand 3D visual features. Comparing the results

of baseline and (a,b,c), we find that using visual embeddings of medical entities substantially enhances the scores, especially when using both the M_E and M_I in KJA (see setting (c)). We speculate the reason is that the explicit and implicit medical relations within M_E and M_I effectively enhance entity visual embeddings with medical grounds, thereby guiding LLM to generate precise texts. Besides, (b) highly improves F1 by 8.5% over (a), which illustrates the explicit entity relations extracted from the expert knowledge bring more positive effects than predicted relations. We also observe that (d) and (e) exhibit significant improvements over (a) and (b), respectively, especially in Recall and CIDEr. This indicates the contribution of learning status, which guides LLM to balance the learning of diverse entities and generate precise findings. Our model’s lower precision compared to (c) stems from the reasons discussed in Section 4.2.

To further illustrate the advantages of our entity-balanced prompt, we compare it with prompts that rely on classified categories (Jin et al. 2024; Zhou et al. 2024). We collect the categories (“exist” or “not exist”) of each entity, which are predicted by our medical entity classifier. These categories are mapped into textual embeddings (denoted as *Category*.) and added to the prompt. Comparing the first two lines and the last two lines in Table 3, it is clear that *Category*. decreases the scores, which is inconsistent with current 2D MRG works. The reason may be that for more complex 3D data, classified category allows LLMs to take more shortcuts than interpreting visual patterns. This attributes the upper limit of performance to the classifier, rather than LLM itself.

Method	Entity-balanced Prompts		Entity Adjacency Matrix		CE Metrics			NLG Metrics			
	Entity Visual Embed.	Status Embed.	w/ M_E	w/ M_I	F1	R	P	B4	M	R-L	C
baseline					33.7	43.1	46.3	10.2	24.0	32.3	12.8
(a)	✓		✓		35.4	43.9	45.8	11.3	24.5	35.2	17.1
(b)	✓			✓	43.9	56.1	47.9	13.1	26.5	36.7	19.0
(c)	✓		✓	✓	44.0	55.9	49.3	13.9	26.8	36.4	21.3
(d)	✓	✓	✓		45.4	60.0	44.6	12.7	28.8	34.5	21.7
(e)	✓	✓		✓	49.4	67.8	46.5	14.2	28.2	35.9	19.6
Ours	✓	✓	✓	✓	51.4	68.8	48.8	14.7	29.2	36.8	22.3

Table 2: Ablation studies of the proposed MEPNet on the BCT-CHR dataset, where different settings of medical entity-balanced prompts and knowledge-driven joint attention are evaluated. KJA denotes the integration of knowledge-driven joint attention. M_E and M_I denote the medical adjacency matrix and predicted adjacency matrix in KJA, respectively.

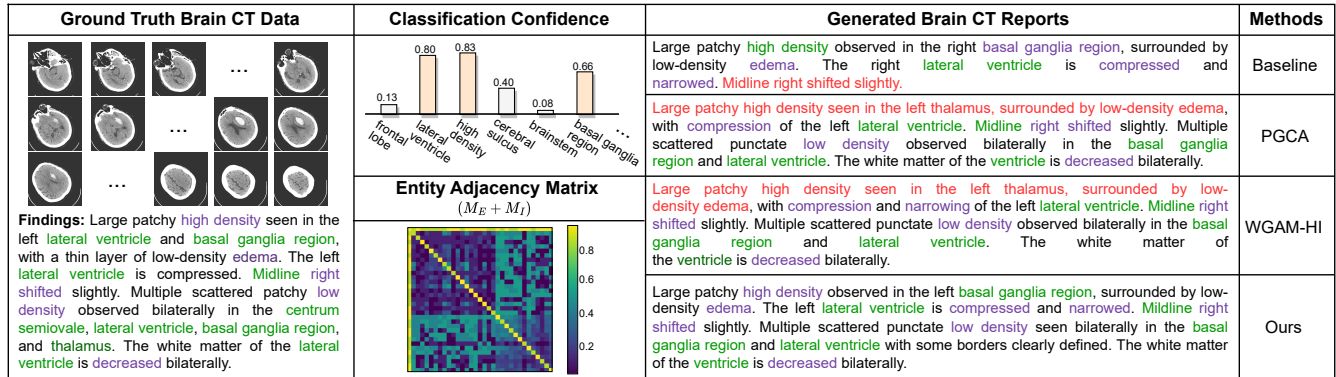


Figure 6: Comparison of brain CT reports generated by different models. Correct medical entities of anatomy region and lesion are marked in green and purple, respectively. Incorrect contents are in red. Classification confidence and entity adjacency matrix are also visualized for supplements.

Entity Prompts			B4	M	R-L	C	F1
Embed.	Status.	Category.					
✓			13.9	26.8	36.4	21.3	44.0
✓		✓	13.3	26.2	36.3	18.9	45.8
✓	✓		14.7	29.2	36.8	22.3	51.4
✓	✓	✓	14.2	29.1	35.4	20.6	46.6

Table 3: Comparisons on using different prompts. *Embed.*, *Status.* and *Category.* denote the entity visual embedding, status embedding, and classified entity category, respectively.

4.4 Qualitative Studies

In Figure 6, we give examples of brain CT reports to better understand our approach. As we can see, the baseline generates the worst report with errors and missed findings. PGCA and WGAM-HI generate more detailed medical content but still have critical errors, such as misdiagnosing the “high density seen in the left thalamus” and “surrounded by low-density edema”, which indicates the weakness in capturing diverse medical entities. In contrast, our model correctly describes a border range of medical entities, e.g. *basal ganglia*, *lateral ventricles*, and *high density*. It suggests that, in our approach,

LLM is supported by enriched medical clues within entity visual embeddings and learning status, thereby aligning well with ground truth reports. To further understand the learning process of medical entities, we visualize the confidence score of the medical entity classifier, which manifests aligned results with entities in ground truth reports. We also visualize the medical entity adjacency matrix (a weighted sum of M_E and M_I calculated by Equation 6). It reveals that the model can integrate both explicit and implicit medical knowledge and leverage it for report generation.

Conclusion

In this paper, we propose to boost brain CT report generation via a medical entity-balanced prompting network, namely MEPNet, which seamlessly enhances the capabilities of the LLM to fairly learn and generate various medical entities. Experiment results show the effectiveness of our model in comparison with SoTA methods on both the BCT-CHR and CTRG-Brain benchmarks. We also verify the improvements brought by each prompt component, i.e. entity visual embedding and learning status embedding. Notably, our approach provides a practical paradigm for activating the capability of LLMs to balance the learning of diverse entities, showing the potential of extending the approach to other medical tasks.

Acknowledgements

This work was supported in part by the National Natural Science Foundation of China under Grant 61906007, 62276010, and 62306253, in part by the Guangdong Natural Science Fund-General Programme under Grant 2024A1515010233.

References

- Anderson, P.; He, X.; Buehler, C.; Teney, D.; Johnson, M.; Gould, S.; and Zhang, L. 2018. Bottom-up and top-down attention for image captioning and visual question answering. In *2018 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2018*, 6077–6086.
- Bhadauria, H. S.; and Dewal, M. L. 2014. Intracranial hemorrhage detection using spatial fuzzy c-mean and region-based active contour on brain CT imaging. *Signal Image Video Process.*, 8(2): 357–364.
- Blankemeier, L.; Cohen, J. P.; Kumar, A.; Veen, D. V.; Gardezi, S. J. S.; Paschali, M.; Chen, Z.; Delbrouck, J.; Reis, E. P.; Truys, C.; Bluethgen, C.; Jensen, M. E. K.; Ostmeier, S.; Varma, M.; Valanarasu, J. M. J.; Fang, Z.; Huo, Z.; Nabulsi, Z.; Ardila, D.; Weng, W.; Júnior, E. A.; Ahuja, N.; Fries, J. A.; Shah, N. H.; Johnston, A.; Boutin, R. D.; Wentland, A.; Langlotz, C. P.; Hom, J.; Gatidis, S.; and Chaudhari, A. S. 2024. Merlin: A Vision Language Foundation Model for 3D Computed Tomography. *CoRR*, abs/2406.06512.
- Brady, A.; Laoide, R. Ó.; McCarthy, P.; and McDermott, R. 2012. Discrepancy and error in radiology: concepts, causes and consequences. *The Ulster medical journal*, 81(1): 3.
- Bu, S.; Song, Y.; Li, T.; and Dai, Z. 2024. Dynamic Knowledge Prompt for Chest X-ray Report Generation. In *Proceedings of the 2024 Joint International Conference on Computational Linguistics, Language Resources and Evaluation, LREC/COLING*, 5425–5436.
- Chen, Z.; Luo, L.; Bie, Y.; and Chen, H. 2024. Dia-LLaMA: Towards Large Language Model-driven CT Report Generation. *CoRR*, abs/2403.16386.
- Chen, Z.; Shen, Y.; Song, Y.; and Wan, X. 2021. Cross-modal Memory Networks for Radiology Report Generation. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing, ACL/IJCNLP 2021*, 5904–5914.
- Chilamkurthy, S.; Ghosh, R.; Tanamala, S.; Biviji, M.; Campeau, N. G.; Venugopal, V. K.; Mahajan, V.; Rao, P.; and Warier, P. 2018. Development and validation of deep learning algorithms for detection of critical findings in head CT scans. *arXiv preprint arXiv:1803.05854*.
- He, J.; Li, P.; Liu, G.; Zhao, Z.; and Zhong, S. 2024. PeFoMed: Parameter Efficient Fine-tuning on Multimodal Large Language Models for Medical Visual Question Answering. *arXiv preprint arXiv:2401.02797*.
- He, K.; Zhang, X.; Ren, S.; and Sun, J. 2016. Deep Residual Learning for Image Recognition. In *2016 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2016*, 770–778.
- Hu, E. J.; Shen, Y.; Wallis, P.; Allen-Zhu, Z.; Li, Y.; Wang, S.; Wang, L.; and Chen, W. 2021. Lora: Low-rank adaptation of large language models. *arXiv preprint arXiv:2106.09685*.
- Ji, J.; Xu, C.; Zhang, X.; Wang, B.; and Song, X. 2020. Spatio-Temporal Memory Attention for Image Captioning. *IEEE Trans. Image Process.*, 29: 7615–7628.
- Jia, X.; Xiong, Y.; Zhang, J.; Zhang, Y.; and Zhu, Y. 2020. Few-shot Radiology Report Generation for Rare Diseases. In *IEEE International Conference on Bioinformatics and Biomedicine, BIBM 2020*, 601–608.
- Jin, H.; Che, H.; Lin, Y.; and Chen, H. 2024. PromptMRG: Diagnosis-Driven Prompts for Medical Report Generation. In *Thirty-Eighth AAAI Conference on Artificial Intelligence, AAAI 2024*, 2607–2615.
- Jing, B.; Xie, P.; and Xing, E. P. 2018. On the Automatic Generation of Medical Imaging Reports. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics, ACL 2018*, 2577–2586.
- Lavie, A.; and Agarwal, A. 2007. METEOR: An Automatic Metric for MT Evaluation with High Levels of Correlation with Human Judgments. In *Proceedings of the Second Workshop on Statistical Machine Translation, WMT@ACL 2007*, 228–231.
- Li, C. Y.; Liang, X.; Hu, Z.; and Xing, E. P. 2019. Knowledge-Driven Encode, Retrieve, Paraphrase for Medical Image Report Generation. In *The Thirty-Third AAAI Conference on Artificial Intelligence, AAAI 2019*, 6666–6673.
- Li, M.; Lin, B.; Chen, Z.; Lin, H.; Liang, X.; and Chang, X. 2023. Dynamic Graph Enhanced Contrastive Learning for Chest X-ray Report Generation. In *IEEE Conference on Computer Vision and Pattern Recognition*, 3334–3343.
- Lin, C.-Y. 2004. Rouge: A package for automatic evaluation of summaries. In *Text summarization branches out*, 74–81.
- Liu, C.; Tian, Y.; Chen, W.; Song, Y.; and Zhang, Y. 2024a. Bootstrapping Large Language Models for Radiology Report Generation. In *Thirty-Eighth AAAI Conference on Artificial Intelligence, AAAI 2024*, 18635–18643.
- Liu, C.; Wan, Z.; Wang, Y.; Shen, H.; Wang, H.; Zheng, K.; Zhang, M.; and Arcucci, R. 2024b. Benchmarking and Boosting Radiology Report Generation for 3D High-Resolution Medical Images. *CoRR*, abs/2406.07146.
- Liu, F.; Wu, X.; Ge, S.; Fan, W.; and Zou, Y. 2021. Exploring and Distilling Posterior and Prior Knowledge for Radiology Report Generation. In *IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2021*, 13753–13762.
- Meta, A. 2024. Introducing meta llama 3: The most capable openly available llm to date. *Meta AI*.
- Papineni, K.; Roukos, S.; Ward, T.; and Zhu, W. 2002. Bleu: a Method for Automatic Evaluation of Machine Translation. In *Proceedings of the 40th Annual Meeting of the Association for Computational Linguistics, ACL 2002*, 311–318.
- Qin, H.; and Song, Y. 2022. Reinforced Cross-modal Alignment for Radiology Report Generation. In *Findings of the Association for Computational Linguistics: ACL 2022*, 448–458.

- Shen, Q.; Shi, Y.; Zhang, X.; Ji, J.; Liu, Y.; and Xu, H. 2024. GHCL: Gaussian heuristic curriculum learning for Brain CT report generation. *Multim. Syst.*, 30(2): 69.
- Shi, Y.; Ji, J.; Zhang, X.; Liu, Y.; Wang, Z.; and Xu, H. 2024. Prior tissue knowledge-driven contrastive learning for brain CT report generation. *Multim. Syst.*, 30(2): 98.
- Shi, Y.; Ji, J.; Zhang, X.; Qu, L.; and Liu, Y. 2023. Granularity Matters: Pathological Graph-driven Cross-modal Alignment for Brain CT Report Generation. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing, EMNLP*, 6617–6630.
- Song, X.; Zhang, X.; Ji, J.; Liu, Y.; and Wei, P. 2022. Cross-modal Contrastive Attention Model for Medical Report Generation. In *Proceedings of the 29th International Conference on Computational Linguistics, COLING 2022*, 2388–2397.
- Tang, Y.; Yang, H.; Zhang, L.; and Yuan, Y. 2024. Work like a doctor: Unifying scan localizer and dynamic generator for automated computed tomography report generation. *Expert Systems with Applications*, 237: 121442.
- Touvron, H.; Martin, L.; Stone, K.; Albert, P.; Almahairi, A.; Babaei, Y.; Bashlykov, N.; Batra, S.; Bhargava, P.; Bhosale, S.; Bikel, D.; Blecher, L.; Canton-Ferrer, C.; Chen, M.; Cucurull, G.; Esiobu, D.; Fernandes, J.; Fu, J.; Fu, W.; Fuller, B.; Gao, C.; Goswami, V.; Goyal, N.; Hartshorn, A.; Housseini, S.; Hou, R.; Inan, H.; Kardas, M.; Kerkez, V.; Khabsa, M.; Kloumann, I.; Korenev, A.; Koura, P. S.; Lachaux, M.; Lavril, T.; Lee, J.; Liskovich, D.; Lu, Y.; Mao, Y.; Martinet, X.; Mihaylov, T.; Mishra, P.; Molybog, I.; Nie, Y.; Poulton, A.; Reizenstein, J.; Rungta, R.; Saladi, K.; Schelten, A.; Silva, R.; Smith, E. M.; Subramanian, R.; Tan, X. E.; Tang, B.; Taylor, R.; Williams, A.; Kuan, J. X.; Xu, P.; Yan, Z.; Zarov, I.; Zhang, Y.; Fan, A.; Kambadur, M.; Narang, S.; Rodriguez, A.; Stojnic, R.; Edunov, S.; and Scialom, T. 2023. Llama 2: Open Foundation and Fine-Tuned Chat Models. *CoRR*, abs/2307.09288.
- Vedantam, R.; Zitnick, C. L.; and Parikh, D. 2015. CIDER: Consensus-based image description evaluation. In *IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2015*, 4566–4575.
- Wang, J.; Bhalerao, A.; and He, Y. 2022. Cross-modal prototype driven network for radiology report generation. In *Computer Vision - ECCV 2022 - 17th European Conference*, 563–579. Springer.
- Wang, X.; Peng, Y.; Lu, L.; Lu, Z.; and Summers, R. M. 2018. Tienet: Text-image embedding network for common thorax disease classification and reporting in chest x-rays. In *2018 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2018*, 9049–9058.
- Xu, K.; Ba, J.; Kiros, R.; Cho, K.; Courville, A. C.; Salakhutdinov, R.; Zemel, R. S.; and Bengio, Y. 2015. Show, attend and tell: Neural image caption generation with visual attention. In *Proceedings of the 32nd International Conference on Machine Learning, ICML 2015*, 2048–2057.
- Yan, A.; He, Z.; Lu, X.; Du, J.; Chang, E. Y.; Gentili, A.; McAuley, J. J.; and Hsu, C. 2021. Weakly Supervised Contrastive Learning for Chest X-Ray Report Generation. In *Findings of the Association for Computational Linguistics: EMNLP 2021*, 4009–4015.
- Yan, B.; Liu, R.; Kuo, D. E.; Adithan, S.; Reis, E. P.; Kwak, S.; Venugopal, V. K.; O’Connell, C.; Saenz, A.; Rajpurkar, P.; and Moor, M. 2023. Style-Aware Radiology Report Generation with RadGraph and Few-Shot Prompting. In *Findings of the Association for Computational Linguistics: EMNLP 2023*, 14676–14688.
- Yang, S.; Ji, J.; Zhang, X.; Liu, Y.; and Wang, Z. 2021. Weakly Guided Hierarchical Encoder-Decoder Network for Brain CT Report Generation. In *IEEE International Conference on Bioinformatics and Biomedicine, BIBM 2021*, 568–573. IEEE.
- Yang, S.; Wu, X.; Ge, S.; Zheng, Z.; Zhou, S. K.; and Xiao, L. 2023. Radiology report generation with a learned knowledge base and multi-modal alignment. *Medical Image Anal.*, 86: 102798.
- Zhang, K.; Jiang, H.; Zhang, J.; Huang, Q.; Fan, J.; Yu, J.; and Han, W. 2023a. Semi-supervised Medical Report Generation via Graph-guided Hybrid Feature Consistency. *IEEE Transactions on Multimedia*.
- Zhang, X.; Yang, S.; Shi, Y.; Ji, J.; Liu, Y.; Wang, Z.; and Xu, H. 2023b. Weakly guided attention model with hierarchical interaction for brain CT report generation. *Computers in Biology and Medicine*, 107650.
- Zhang, Y.; Wang, X.; Xu, Z.; Yu, Q.; Yuille, A. L.; and Xu, D. 2020. When Radiology Report Generation Meets Knowledge Graph. In *The Thirty-Fourth AAAI Conference on Artificial Intelligence, AAAI 2020*, 12910–12917.
- Zheng, C.; Ji, J.; Shi, Y.; Zhang, X.; and Qu, L. 2024. See Detail Say Clear: Towards Brain CT Report Generation via Pathological Clue-driven Representation Learning. In *Findings of the Association for Computational Linguistics: EMNLP 2024*, 16542–16552.
- Zhou, Z.; Shi, M.; Wei, M.; Alabi, O.; Yue, Z.; and Vercauteren, T. 2024. Large Model driven Radiology Report Generation with Clinical Quality Reinforcement Learning. *CoRR*, abs/2403.06728.