

DAMPER: A Dual-Stage Medical Report Generation Framework with Coarse-Grained MeSH Alignment and Fine-Grained Hypergraph Matching

Xiaofei Huang¹, Wenting Chen^{2*}, Jie Liu², Qisheng Lu¹, Xiaoling Luo¹, Linlin Shen^{1,3,4*}

¹Computer Vision Institute, College of Computer Science and Software Engineering, Shenzhen University, Shenzhen, China

²Department of Electrical Engineering, City University of Hong Kong, Kowloon, Hong Kong

³National Engineering Laboratory for Big Data System Computing Technology, Shenzhen University

⁴Guangdong Provincial Key Laboratory of Intelligent Information Processing

2450101014@mails.szu.edu.cn, {wentichen7-c, jilu.ee}@my.cityu.edu.hk, qishenglu85@gmail.com, {lluo, llshen}@szu.edu.cn

Abstract

Medical report generation is crucial for clinical diagnosis and patient management, summarizing diagnoses and recommendations based on medical imaging. However, existing work often overlook the clinical pipeline involved in report writing, where physicians typically conduct an initial quick review followed by a detailed examination. Moreover, current alignment methods may lead to misaligned relationships. To address these issues, we propose DAMPER, a dual-stage framework for medical report generation that mimics the clinical pipeline of report writing in two stages. In the first stage, a MeSH-Guided Coarse-Grained Alignment (MCG) stage that aligns chest X-ray (CXR) image features with medical subject headings (MeSH) features to generate a rough keyphrase representation of the overall impression. In the second stage, a Hypergraph-Enhanced Fine-Grained Alignment (HFG) stage that constructs hypergraphs for image patches and report annotations, modeling high-order relationships within each modality and performing hypergraph matching to capture semantic correlations between image regions and textual phrases. Finally, the coarse-grained visual features, generated MeSH representations, and visual hypergraph features are fed into a report decoder to produce the final medical report. Extensive experiments on public datasets demonstrate the effectiveness of DAMPER in generating comprehensive and accurate medical reports, outperforming state-of-the-art methods across various evaluation metrics.

Introduction

Medical reports are essential for highlighting clinical findings and summarizing diagnoses and recommendations, which are crucial for clinical decision-making and patient management. However, the rise in medical imaging has made traditional manual report generation time-consuming and error-prone, potentially affecting patient safety. Therefore, automated medical report generation techniques have become increasingly important (Cao et al. 2023).

With the widespread application of deep learning in the medical field, advancements have also been made in various tasks such as medical visual question answering (Huang and Gong 2024), rectal cancer diagnosis (Zhang et al. 2024), clinical decision making (Liu et al. 2024a) and retinal disease detection (Luo et al. 2023, 2024; Xu et al. 2024). Au-

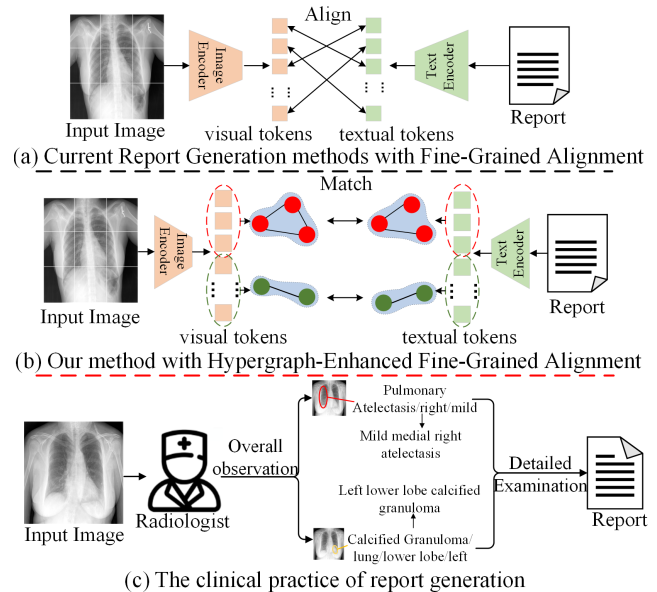


Figure 1: (a) Represents the previous research works on fine-grained alignment. (b) Our research work on fine-grained alignment in this paper. (c) The motivation for our approach: Mimicking the real process of physicians writing reports.

tomatic algorithms for medical report generation have been widely studied. Some methods (Shin et al. 2016; Yuan et al. 2019) directly map input chest X-ray (CXR) images to medical reports using CNNs and RNNs, while others focus on effective image-text alignment to facilitate report generation. The latter methods believe that cross-modal disparity between images and texts may affect the mapping from CXR images to reports. They primarily rely on two types of alignment: coarse-grained and fine-grained alignment. Coarse-grained alignment methods (Yang et al. 2023; Chen et al. 2024b,c) align entire image features with whole report features, whereas fine-grained alignment methods (Qin and Song 2022; Yang et al. 2023) aim to match specific image regions with corresponding descriptions in the report. However, these methods still face two main challenges when using alignment for report generation.

*Corresponding Authors

The first challenge in automatic medical report generation is that current methods often generate reports directly from medical images, overlooking the clinical pipeline involved in report writing. Physicians typically conduct a quick review of images to form an overall impression, identifying major structures and abnormalities, followed by a detailed examination of specific regions for accurate diagnosis (Brady 2018), as shown in Fig. 1 (c). While some approaches (Mei et al. 2024; Jing, Xie, and Xing 2017) consider the clinical practice of writing reports, they rely on clinical history or use disease labels to generate key phrases. However, these methods require additional effort to obtain extra patient information, which may not always be readily available due to privacy concerns surrounding medical data. Hence, it is imperative to consider the clinical practice of report generation to ensure the comprehension and accuracy of the generated reports.

Another challenge in current fine-grained alignment methods is that they directly match image tokens with words, leading to misaligned relationships. In medical reports, detailed pathological findings are often described for different anatomical regions (Tanida et al. 2023), as shown in Fig. 1 (c), "left lower lobe calcified granuloma" (yellow region) and "mild right lung atelectasis" (red region). This implies that certain phrases in the report are correlated with the specific anatomical regions. Existing methods (Zhao et al. 2023; Li et al. 2023) achieve fine-grained alignment by establishing the relation between each image region and each word token one by one, as shown in Fig. 1 (a). However, this direct token-to-word matching approach can potentially fragment semantic information across different patches or words, resulting in incomplete and unreasonable alignments between individual patches and words. Therefore, it is crucial to consider intra-modal relationships during the fine-grained alignment process in report generation.

To address these challenges, we propose the **Dual-Stage Medical Report Generation Framework with Coarse-Grained MeSH Alignment and Fine-Grained Hypergraph Matching (DAMPER)**. This framework is inspired by the actual process of medical report writing and aims to incorporate clinical practices. DAMPER consists of a MeSH-guided Coarse-Grained Alignment (MCG) to mimic radiologists' overall observations and a Hypergraph-Enhanced Fine-Grained Alignment (HFG) to replicate detailed examinations. We found that MeSH information extracted from medical reports effectively captures key content and details, aiding in simulating the report writing process in clinical practice. **MCG stage** utilizes MeSH information to represent the patient's overall health status and align CXR images with MeSH. It consists of a MeSH Encoding module that transforms MeSH into graph embeddings, a GAN-Based MeSH-CXR Alignment (MCA) module that aligns multi-view CXR images with these embeddings, and a CXR-to-MeSH generation (CMG) module that maps the aligned CXR features into MeSH. The MCA module employs adversarial learning to ensure that the multi-view features generated are close to the graph embeddings. To further enhance overall observation, the CMG uses attention

fusion and a MeSH Decoder to convert multi-view features into MeSH. The hypergraph structure effectively captures complex intra-modal relationships during fine-grained alignment. **HFG stage**, we address potential misalignments between image patches and words by introducing hypergraphs to consider intra-modal relationships, preserving higher-order relationships throughout the alignment process. HFG includes Intra-Patient Report-CXR Alignment (Intra-RCA), Inter-Patient Report-CXR Alignment (Inter-RCA), and a report decoder. Intra-RCA constructs hypergraphs and performs node matching. We use hyperedges to connect patches within image regions and words within report phrases, capturing the complex relationships among modality-specific elements. Through node matching, we effectively align image patches with their corresponding report phrases. Inter-RCA employs contrastive learning for hypergraph embedding matching, optimizing feature representations to enhance accuracy in matching image regions with report phrases. After attention-based fusion of outputs from MCA, CMG, and Intra-RCA, the results are fed into the report decoder to generate a comprehensive medical report. To manage varying numbers of input images, we introduce a Bernoulli indicator in the generative adversarial and graph matching processes, ensuring model robustness even with missing visual information. Our contributions can be summarized as follows:

- We introduce DAMPER, a medical report generation model designed with a dual-stage architecture that integrates coarse-grained and fine-grained alignments. This model emulates the process physicians follow when creating reports, enabling the generation of comprehensive and accurate medical documentation.
- To simulate the radiologist's holistic observation, we designed the MCG stage, which aligns CXR visual features with MeSH information.
- To simulate the detailed examination following the initial overall observation, we developed the HFG stage, aligning image regions with report phrases while preserving the high-order relationships among image patches.
- Experimental evaluations of DAMPER were conducted on two public datasets, IU-Xray and MIMIC-CXR. The results demonstrate that DAMPER outperforms state-of-the-art methods across various evaluation metrics.

Related Work

Cross-Modality Alignment

Cross-modal alignment involves matching and integrating features from different modalities, such as images and text, to effectively combine visual and semantic information. Recent advancements in this field include techniques like CLIP, GAN, and graph matching.

The CLIP model excels in coarse-grained feature interaction through pretraining on large image-text datasets, demonstrating strong zero-shot learning capabilities (Radford et al. 2021). CLIP's success has spurred further research, such as SoftCLIP (Gao et al. 2024) and mCLIP (Chen et al. 2023), which have refined alignment techniques.

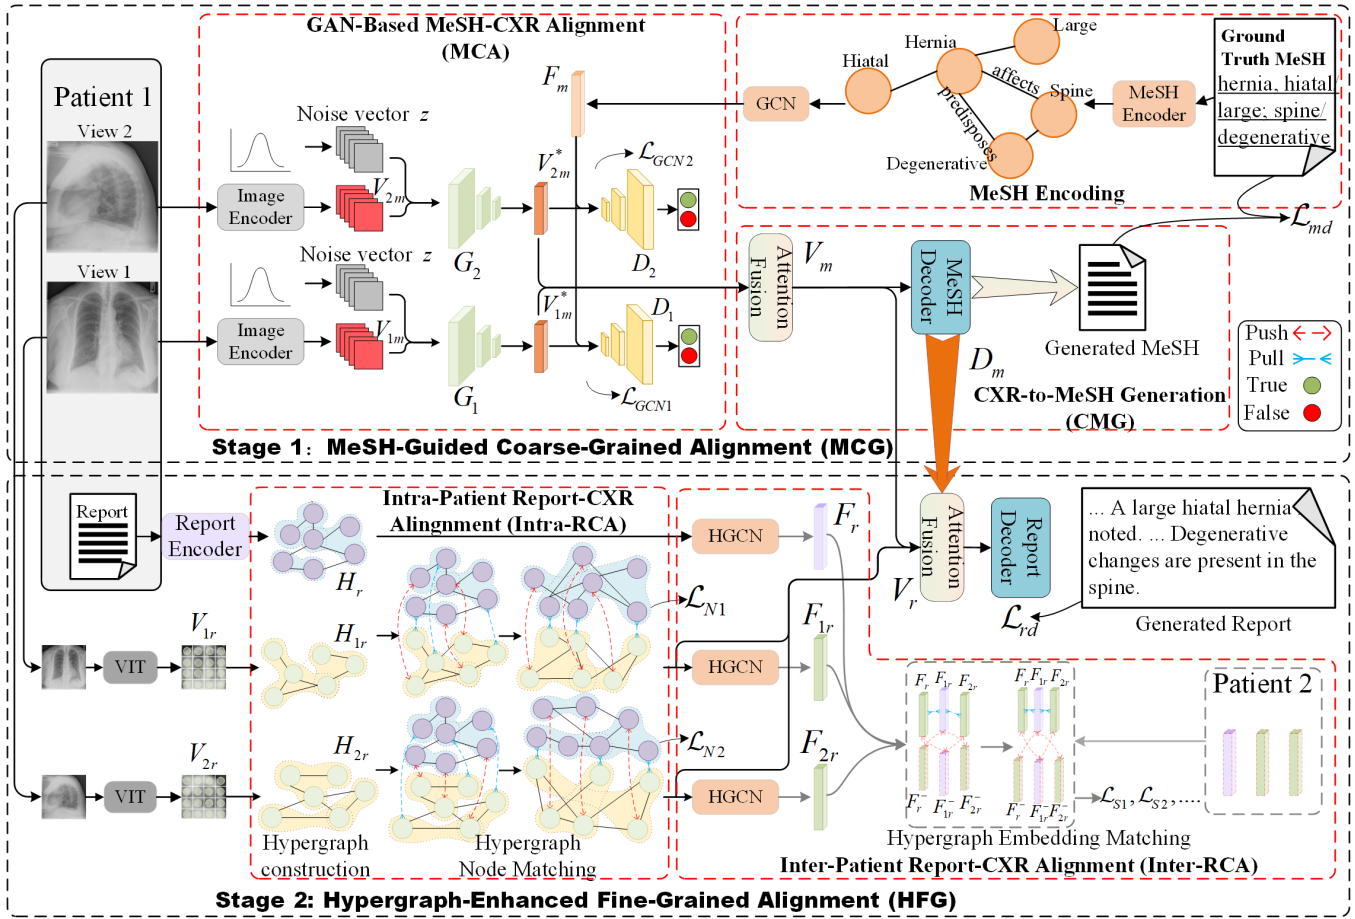


Figure 2: Overview of the proposed DAMPER framework. In the MCG stage, the MCA module extracts coarse-grained visual features aligned with MeSH terms, which are then input into the CMG. The HFG stage includes Intra-RCA and Inter-RCA modules that acquire detailed information corresponding to the report. This information, along with the output from MCA and CMG, is fed into the report decoder to generate the final medical report.

However, since MeSH consists of keywords, it may not fully capture the semantic relationships between images, and CLIP-like models often require extensive training data for specialized domains. In contrast, GAN offers greater flexibility, making it more suitable for the coarse-grained alignment task in this context.

Graph matching techniques, such as GEM (Liu et al. 2024b) and Bi-VLGM (Chen et al. 2024a), provide a novel approach to managing complex cross-modal relationships. These methods utilize graph structures to represent intricate associations between image regions and text, allowing for more precise alignment. Our study builds on this by exploring hypergraph matching for fine-grained alignment, which enhances the representation of complex relationships between visual features and text phrases.

Medical Report Generation

Medical report generation is a complex task that involves transforming medical images into accurate and coherent textual descriptions. To improve the quality of generated text, researchers introduced Transformer-based methods (Yang

et al. 2023; Cao et al. 2022), significantly enhancing the coherence and contextual understanding of generated reports. Subsequently, some researchers attempted to incorporate additional information into the decoding process to further improve the quality and accuracy of the generated reports (Zhang et al. 2021; Liu et al. 2021).

Despite the significant advancements in report generation quality and medical accuracy achieved by these methods, they share a common limitation: the MeSH is not fully explored. MeSH terms are crucial elements that encapsulate disease conditions and health status, and they are essential for generating reports that accurately reflect pathological information. In light of this, we propose incorporating MeSH information into the report generation process to better simulate the real-world report writing workflow, with the aim of producing more accurate radiology reports.

Methods

Overview

DAMPER is a two-stage framework that simulates the medical report generation process in clinical practice. The first stage, MCG, includes a MeSH Encoding, MCA, and CMG modules. During the MCG stage, we extract MeSH graph embeddings F_m and image representations V_{im} using the MeSH Encoding and Image Encoder. Then, the MCA module employs a GAN to generate coarse-grained visual features V_{im}^* aligned with F_m , which are input into the CMG module to produce phrase representations that capture the overall impression. The second stage, HFG, involves features extraction, Intra-RCA, Inter-RCA, and report generation. HFG constructs hypergraphs H_r and H_{ir} based on report phrases and image regions, achieving precise matching through loss functions \mathcal{L}_N and \mathcal{L}_S in Intra-RCA and Inter-RCA. Finally, the outputs V_m , D_m , and V_r from MCA, CMG, and Intra-RCA are fused and passed into the report decoder to generate the textual report. The overall workflow is illustrated in Fig. 2.

Stage 1: MeSH-Guided Coarse-Grained Alignment

Brady (2018) noted that in clinical report writing, physicians first form an overall impression before analyzing details. Existing methods either require more patient information (Mei et al. 2024) or overlook underlying pathologies when considering clinical practices (Jing, Xie, and Xing 2017). To address this, we designed the MCG stage, which includes the MeSH Encoding, MCA, and CMG. The primary objective of MCG is to optimize the extraction of coarse-grained visual information and align it with MeSH, simulating the radiologist’s overall assessment to accurately capture pathological information and guide the subsequent report generation.

MeSH Encoding We use Clinical-BERT (Yan and Pei 2022) to extract word information from MeSH as node features N_m . Reference to RGM (Fu et al. 2021), we obtain the corresponding edge relationship features E_m to construct the MeSH graph $G_m = (N_m, E_m)$. This graph is then input into a GCN to obtain its embedding representation:

$$F_m = GCN(G_m). \quad (1)$$

GAN-Based MeSH-CXR Alignment (MCA) Given the flexibility of GANs and the need to obtain visual features corresponding to MeSH, we developed the MCA module. In practical applications, multiple medical images, denoted as I_i for $i = 1, 2, \dots, H$, may need to be processed simultaneously. For each image I_i , we use ResNet to extract its overall visual feature V_{im} . During the adversarial learning, we use the generator of GAN to generate feature vectors V_{im}^* that resembles the target vector F_m based on the initial feature vector V_{im} . The loss for this adversarial generation process for view i is expressed as:

$$\mathcal{L}_{GANi} = \mathbb{E}_{x \sim p_{data}(x)} [\log D_i(x|V_{im})] + \mathbb{E}_{z \sim p_z(z)} [\log(1 - D_i(G_i(z, V_{im})|V_{im}))] + \|F_m - G_i(z, V_{im})\|^2. \quad (2)$$

The loss function consists of three components: the discriminator’s loss on real data, the discriminator’s loss on

generated data, and the mean squared error between the generated features and the MeSH features. Here, D_i and G_i represent the discriminator and generator in the GAN for view i , respectively, while z is the noise vector generated from a normal distribution.

CXR-to-MeSH Generation (CMG) To enhance the holistic representation and provide more accurate feature representations for report generation, we designed the CMG module. It employs an attention mechanism to fuse the multimodal information V_{-m}^* , obtaining a comprehensive representation from all viewpoints, can be expressed as:

$$V_m = PA(Q, K, V) = \text{Softmax}\left(\frac{QK^T}{\sqrt{d_k}}\right)V, \quad (3)$$

$$Q = XW_Q, K = XW_K, V = XW_V, \quad (4)$$

$$X = [\delta^M \cdot V_{1m}^*; \delta^M \cdot V_{2m}^*; \dots; \delta^M \cdot V_{Hm}^*], \quad (5)$$

where W_Q , W_K , and W_V are learnable parameter matrices used to generate queries, keys, and values, respectively. $\delta \in \{0, 1\}$ is a Bernoulli indicator introduced to enhance the model’s robustness in scenarios where certain views are missing. In this manner, we obtain the coarse-grained radiological representation V_m after the MCA process. The V_m is fed into the MeSH decoder to generate the corresponding MeSH representation, which we train using a negative log-likelihood loss function:

$$\mathcal{L}_{md} = -\sum_{t=1}^T \log P(y_t|y_{<t}, V_m), \quad (6)$$

here, y_t represents the t -th word in the target sequence, and $y_{<t}$ represents the preceding words. To further distill the essential MeSH information, we apply a SoftPool operation on the final layer’s hidden features, obtaining a refined MeSH feature representation.

The loss function for the entire MCG stage can be expressed as follows:

$$\mathcal{L}_{MCG} = \sum_{i=1}^H \delta^M \mathcal{L}_{GANi} + \mathcal{L}_{md}. \quad (7)$$

Stage 2: Hypergraph-Enhanced Fine-Grained Alignment

Previous fine-grained alignment methods directly matched image regions with words (Zhao et al. 2023; Li et al. 2023), often neglecting higher-order semantic relationships within each modality. To achieve precise alignment while preserving complex intra-modal connections, we designed the HFG stage, which includes: Features Extraction, Intra-RCA and Inter-RCA, and Report Generation. The HFG stage aims to refine the extracted fine-grained visual information to match detailed medical expressions in the report, simulating the thorough examination process in clinical report writing.

Features Extraction We utilize Clinical-BERT to extract textual information from medical report: $R = \text{Clinical} - \text{BERT}(\text{report})$. For a given medical image I_i , we employ a Vision Transformer (ViT) model to extract its fine-grained features: $V_{ir} = \text{ViT}(I_i)$.

Intra-Patient Report-CXR Alignment (Intra-RCA) To capture intra-modal relationships and align image regions with report phrases, we designed the Intra-RCA module. We use RGM and KNN clustering to build the report hypergraph $H_r = (N_r, E_r)$, which is input into a HGCN to obtain the overall relational representation $F_r = HGCN(H_r)$. Based on V_{ir} , we construct a visual feature hypergraph $H_{ir} = (N_{ir}, E_{ir})$ for each view i and process it using a HGCN to obtain the overall representation $F_{ir} = HGCN(H_{ir})$.

We employ node feature matching to reduce the distance between fine-grained visual features and similar textual expressions while increasing the disparity with dissimilar expressions. The node matching loss is defined as follows:

$$\mathcal{L}_{Ni}(N_r, N_{ir}) = \max[\tau - m_{node}(N_r, N_{ir}) + m_{node}(N_r^-, N_{ir}), 0] + \max[\tau - m_{node}(N_r, N_{ir}) + m_{node}(N_r, N_{ir}^-), 0], \quad (8)$$

where τ is the margin parameter, N_r^- and N_{ir}^- denote the negative samples in N_r and N_{ir} , respectively. The graph node matching function $m_{node}(\cdot)$ is used to compute the sum of the maximum similarity values for each node match, which is specifically calculated as follows:

$$m_{node}(N_I, N_T) = \sum_{0 \leq i \leq m, 0 \leq j \leq l} \max[\cos(N_{Ii}, N_{Tj}), 0], \quad (9)$$

where $\cos(\cdot)$ denotes the cosine similarity function.

Inter-Patient Report-CXR Alignment (Inter-RCA) We designed the Inter-RCA module, which employs contrastive learning to achieve hypergraph embedding alignment between positive samples and the hardest negative samples, further enhancing the matching accuracy between image regions and report phrases. We utilize hypergraph embeddings F_r and F_{ir} to represent the features of the entire graphs H_r and H_{ir} , respectively, encompassing both node features and relational characteristics. We introduce a hypergraph embedding matching loss to maximize the distance between negative and positive samples within a mini-batch. The loss function is defined as:

$$\mathcal{L}_{Si}(F_r, F_{ir}) = \max[\gamma - \cos(F_r, F_{ir}) + \cos(F_r^-, F_{ir}), 0] + \max[\gamma - \cos(F_r, F_{ir}) + \cos(F_r, F_{ir}^-), 0], \quad (10)$$

where γ is the margin parameter. F_r^- and F_{ir}^- represent the most challenging negative samples for F_r and F_{ir} , respectively, within the mini-batch.

We concatenate the obtained V_{ir} to derive the fine-grained feature information V_r from all views.

Report Generation The report decoder generates a complete medical report by integrating V_m , V_r , and D_m as inputs using an attention mechanism. By combining these multi-granularity features, it produces a comprehensive report that includes accurate medical subject headings and detailed descriptions. We train the report decoder using a negative log-likelihood loss:

$$\mathcal{L}_{rd} = - \sum_{t=1}^T \log P(y_t | y_{<t}, D_m, V_m, V_r). \quad (11)$$

The loss in the HFG stage can be expressed as:

$$\mathcal{L}_{HFG} = \sum_{i=1}^H \delta^R \cdot (\mathcal{L}_{Ni}(N_r, N_{ir}) + \mathcal{L}_{Si}(F_r, F_{ir})) + \mathcal{L}_{rd}. \quad (12)$$

The overall loss function of the DAMPER framework combines the losses from the MCG and HFG components, denoted as:

$$\mathcal{L}_{total} = \mathcal{L}_{MCG} + \mathcal{L}_{HFG}. \quad (13)$$

By leveraging this diversified loss function, DAMPER effectively transforms multi-granularity visual details into accurate and comprehensive medical reports.

Experiment

Experimental setup

Datasets We conducted experiments on two publicly available datasets: IU-Xray (Demner-Fushman et al. 2016), MIMIC-CXR (Johnson et al. 2019). **IU-Xray** comprises 7,470 pairs of CXR images and corresponding diagnostic reports, each including MeSH information. We used the same data preprocessing and splitting approach as R2Gen, with a training/validation/testing ratio of 7:1:2. **MIMIC-CXR** is the largest publicly available radiology dataset, containing 377,110 CXR images and 227,835 reports from 65,379 patients. We employed the official data split with a training/validation/testing ratio of 7:1:2. Notably, MIMIC-CXR does not include MeSH, so we used the CheXpert (Smit et al. 2020) labeler to generate MeSH based on the reports.

Implementation Details We conducted our experiments using the PyTorch framework on a single NVIDIA V100 GPU. Both the MeSH decoder and the report decoder utilize a 3-layer Transformer decoder architecture. The model was trained using the Adam optimizer with a batch size of 16 and a learning rate of 1e-4.

We compare our method against the following models: AlignTransformer (You et al. 2021), R2Gen (Chen et al. 2020), R2GenCMN (Chen et al. 2021), PPKED (Liu et al. 2021), KiUT (Huang, Zhang, and Zhang 2023), MMTN (Cao et al. 2023), COMG (Gu et al. 2024), MedM2G (Zhan et al. 2024), PromptMRG (Jin et al. 2024), M2KT (Yang et al. 2023), CvT2DistilGPT2 (Nicolson, Dowling, and Koopman 2023), RECAP (Hou et al. 2023a), ICON (Hou et al. 2024) and ORGAN (Hou et al. 2023b). These models represent a diverse range of architectures and methodologies in the field of medical report generation, spanning from deep learning-based encoder-decoder frameworks to multimodal fusion techniques, showcasing the variety and innovation in current research. Through comparisons with these state-of-the-art models, we aim to comprehensively evaluate the performance and advantages of DAMPER in generating high-quality medical reports.

Evaluation Metrics In the evaluation process, we used metrics: BLEU (B) (Papineni et al. 2002), METEOR (M) (Denkowski and Lavie 2011), and ROUGE-L (R-L) (Lin 2004). The BLEU metric can reflect whether the generated reports use correct medical terminology and common expressions. METEOR better reflects the semantic accuracy of the generated reports, rather than just literal matching, helping to assess whether the reports accurately convey medical concepts. ROUGE-L aids in evaluating whether the reports describe various medical findings and diagnoses in the correct order.

Methods	IU-Xray						MIMICI-CXR					
	B-1	B-2	B-3	B-4	M	R-L	B-1	B-2	B-3	B-4	M	R-L
AlignTransformer (MICCAI'21)	0.484	0.313	0.225	0.173	0.204	0.379	0.378	0.235	0.156	0.112	0.158	0.283
R2Gen (EMNLP'20)	0.470	0.304	0.219	0.165	0.187	0.371	0.353	0.218	0.145	0.103	0.142	0.277
R2GenCMN (ACL'21)	0.470	0.304	0.222	0.170	0.191	0.358	0.348	0.206	0.135	0.094	0.136	0.266
PPKD (CVPR'21)	0.483	0.315	0.224	0.168	0.190	0.376	0.360	0.224	0.149	0.106	0.149	0.284
KiUT (CVPR'23)	0.525	0.360	0.251	0.185	0.242	0.409	0.393	0.243	0.159	0.113	0.160	0.285
MMTN (AAAI'23)	0.486	0.321	0.232	0.175	-	0.375	0.379	0.238	0.159	0.116	0.161	0.283
COMG (WACV'24)	0.536	0.378	0.275	0.206	0.218	0.383	0.363	0.235	0.167	0.124	0.128	0.290
MedM2G (CVPR'24)	0.533	0.369	0.278	0.212	-	0.416	0.412	0.260	0.179	0.142	-	0.309
DAMPER	0.520	0.383	0.300	0.225	0.284	0.397	0.402	0.284	0.227	0.193	0.289	0.301

Table 1: Performance of DAMPER and Related Models on IU-Xray and MIMIC-CXR Test Sets.


Input Images	MeSH	Ground-truth report	DAMPER Generated MeSH	DAMPER Generated report	R2Gen Generated report	R2GenCMN Generated report	COMG Generated report
	<p>Nodule/ lung/upper lobe/right ; Osteophyt e/spine</p>	<p>well circumscribed 11 mm right upper lobe nodule . unchanged appearance from previous examination 7 years ago . the trachea is midline . negative for pneumothorax . pleural effusion . the heart size is normal . redemonstrated syndesmophyte .</p>	<p>Nodule/ upper lobe/ right ; Osteophyte/ spine</p>	<p>right upper lobe nodule circumscribed 11 mm . unchanged appearance from previous examination . trachea central . no pneumothorax or pleural effusion . heart size normal . spine bone growth still there .</p>	<p>well-defined 11 mm nodule . its appearance has remained unchanged since the examination 7 years ago . the trachea is in the midline . no pneumothorax or pleural effusion . the heart size is normal . syndesmophyte .</p>	<p>a well-defined mm mass with no change in appearance the examination 7 years ago . the airway is centrally located . no evidence of air or fluid in the chest cavity . the heart size is within normal limits . noted bony growth is still present .</p>	<p>well-defined 11 mm mass with no change in appearance from the examination 7 years ago . the airway is centrally located . no evidence of air or fluid in the chest cavity . the size of the heart is within normal limits . previously noted bony growth is still present .</p>

Figure 3: Visualization of report generation examples includes the input image in the first column, followed by the corresponding MeSH and report in the 2nd and 3rd columns. The 4th column shows the MeSH generated by DAMPER, while the 5th column backward shows reports from various models. MeSH are highlighted in red or green, with related sentences in the generated reports marked accordingly.

Methods	Precision	Recall	F1
R2Gen (EMNLP'20)	0.333	0.273	0.276
R2GenCMN (ACL'21)	0.334	0.275	0.278
CvT2DistilGPT2 (AIM'23)	0.367	0.418	0.391
KiUT (CVPR'23)	0.371	0.318	0.321
M2KT (MIA'23)	0.420	0.339	0.352
COMG (WACV'24)	0.424	0.291	0.345
PromptMRG (AAAI'24)	0.501	0.509	0.476
DAMPER	0.512	0.473	0.507

Table 2: Comparison of CE metrics between DAMPER and related models on the MIMIC-CXR dataset.

To further evaluate the model's clinical performance, we incorporated clinical efficacy (CE) metrics, including precision, recall, and F1-score, following the same approach as (Nicolson, Dowling, and Koopman 2023) and (Yang et al. 2023). On the MIMIC-CXR dataset, we utilized the rule-based CheXpert labeler to extract disease labels from both ground-truth and generated reports, enabling the calculation of CE metrics.

Comparison With State-of-the-Art

In the comparative analysis of the DAMPER model with state-of-the-art approaches on the IU-Xray and MIMIC-

CXR datasets, Table 1 highlights the significant advantages of DAMPER. Compared to other advanced models, DAMPER demonstrated superior performance across both datasets, achieving the highest scores particularly in BLEU-(2-4) and METEOR metrics. Although its performance on BLEU-1 and ROUGE-L was slightly lower than the best-performing models, the substantial improvement in METEOR is especially noteworthy. These results underscore the significant progress made by DAMPER in generating high-quality, semantically accurate medical reports.

Furthermore, Table 2 presents the CE comparison between DAMPER and related models on the MIMIC-CXR dataset, revealing that DAMPER outperformed previous approaches in terms of both precision and F1-score. This demonstrates DAMPER's ability to provide more reliable support for clinical report generation.

A qualitative comparison of DAMPER with open-source models such as R2Gen, R2GenCMN, and COMG is presented, focusing on the accuracy of MeSH terms in the generated reports, as illustrated in Fig. 3. While the reports may appear similar, DAMPER provides more precise descriptions of health conditions and pathological details critical for diagnosis and treatment. For instance, DAMPER accurately captures critical details such as "well circumscribed 11 mm right upper lobe nodule" and "redemonstrated syndesmophyte," whereas other models often give incomplete or inaccurate descriptions. This demonstrates DAMPER's

Status	MCA	CMG	Intra-RCA	Inter-RCA	B-1	B-2	B-3	B-4	M	R-L
w/o CMG+ Intra-RCA + Inter-RCA	✓				0.467	0.313	0.210	0.155	0.188	0.371
w/o Intra-RCA + Inter-RCA	✓	✓			0.501	0.347	0.221	0.163	0.230	0.382
w/o Inter-RCA	✓	✓	✓		0.511	0.358	0.260	0.198	0.260	0.368
w/o Intra-RCA	✓	✓		✓	0.507	0.350	0.258	0.188	0.258	0.353
w/o CMG	✓		✓	✓	0.497	0.341	0.216	0.156	0.193	0.370
w/o MCA		✓	✓	✓	0.478	0.328	0.214	0.153	0.186	0.339
w/o MCG			✓	✓	0.471	0.328	0.215	0.152	0.203	0.348
DAMPER	✓	✓	✓	✓	0.520	0.383	0.300	0.225	0.284	0.397

Table 3: Description of ablation conditions and experimental results on the IU-Xray dataset. A ✓ indicates inclusion.

advantage in integrating MeSH information to produce more comprehensive reports reflecting the patient’s health status. More examples can be found in the appendix.

Ablation Studies

To validate the effectiveness and rationale of each component, we designed the following model configurations (see Table 3) and conducted ablation experiments on the IU-Xray dataset. The results are presented in Table 3.

The Effectiveness of MCA Comparing the test results of w/o MCA with DAMPER reveals a significant performance decline on the IU-Xray dataset when the MCA module is removed, with the METEOR metric dropping by nearly 0.1. Additionally, when comparing w/o CMG with w/o MCG, w/o CMG outperforms w/o MCG in all metrics, with the most notable difference in the ROUGE-L. These results confirm that the MCA module, by focusing on the overall features of the images, significantly enhances the model’s performance and the semantic quality of the generated reports.

The Effectiveness of CMG When comparing the performance of DAMPER with w/o CMG, we observed a significant drop in performance after removing the CMG, with all metrics decreasing by at least 0.03 and METEOR dropping by 0.09. Further analysis of configurations w/o Intra-RCA + Inter-RCA and w/o MCA also showed a decline in performance across all metrics when CMG was removed, though the differences were less pronounced. These findings highlight the importance of CMG in medical report generation, as it enhances semantic accuracy and content consistency, improving the model’s understanding of pathology and health status, and thereby the quality and accuracy of generated reports. Examples of the generated reports can be found in the appendix.

The Effectiveness of Intra-RCA By comparing the results of DAMPER with w/o Intra-RCA and w/o Inter-RCA with w/o Intra-RCA + Inter-RCA, we observed a significant performance drop when Intra-RCA was removed. This decline is likely due to the difficulty in effectively aligning phrases in the report with visual features without the supervision of Intra-RCA. This finding highlights the critical role of Intra-RCA in generating accurate phrase-level text.

The Effectiveness of Inter-RCA Comparing DAMPER with w/o Inter-RCA and w/o Intra-RCA with w/o Intra-

RCA + Inter-RCA, we observed a decline in model performance. The decrease was more pronounced when Intra-RCA was removed, indicating that Intra-RCA is more critical to DAMPER’s performance than Inter-RCA.

Methods	B-1	B-2	B-3	B-4	M	R-L
Disease Labels as MeSH	0.486	0.307	0.213	0.157	0.187	0.365
Ground-truth MeSH (DAMPER)	0.520	0.383	0.300	0.225	0.284	0.397

Table 4: The effectiveness of Ground-truth MeSH information for report generation on the IU-Xray dataset.

The Effectiveness of MeSH Information To evaluate the importance of MeSH information in report generation, we replaced the original MeSH terms in the IU-Xray dataset with those generated by Chexpert, referred to as ”Disease Labels as MeSH.” The results, shown in Table 4, indicate that this replacement leads to a decrease in model performance. However, compared to the results in Table 1, where MCG was omitted, the performance with generated MeSH, while inferior to DAMPER, still surpasses the case where MeSH is entirely discarded. This suggests that although generated MeSH is less effective than the original, it still contributes positively to model performance.

Methods	B-1	B-2	B-3	B-4	M	R-L
w/o Hypergraph + w/ Graph	0.492	0.343	0.220	0.158	0.204	0.369
DAMPER	0.520	0.383	0.300	0.225	0.284	0.397

Table 5: Test Comparison Results for the Validation of Hypergraph effectiveness in DAMPER on the IU-Xray Dataset.

The Effectiveness of Hypergraph In this study, we utilized hypergraphs to establish correspondences between visual region features and report phrases. To assess the superiority of hypergraphs over traditional graphs in this context, we replaced the hypergraphs in DAMPER with graphs, defining this condition as w/o Hypergraph + w/ Graph. The comparative results are presented in Table 5.

The results indicate that replacing the original hypergraphs with graphs leads to a decline in model performance, particularly noticeable in the METEOR metric. This highlights the distinct advantage of hypergraphs in capturing medical terminology and effectively representing the complex relationships between visual features and text. Thus,

hypergraphs are better suited for modeling the intricate associations between visual regions described in reports and the corresponding medical concepts.

Visualization of Coarse-and-Fine-Grained Alignment

We visualized the alignment examples of MCG and HFG in Fig. 4 and Fig. 5 to demonstrate the effectiveness of MCA and hypergraph matching.

Visualization Coarse-Grained Alignment Fig. 4 shows the t-SNE visualization of visual features aligned with MeSH information through GAN. Through visual analysis, we observed a high degree of overlap between the two sets of visual features and MeSH data in the low-dimensional space. This result indicates that MCA successfully aligns MeSH with visual features.

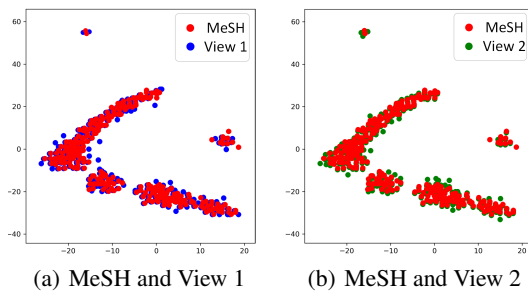


Figure 4: The t-SNE Visualization of Visual Features and MeSH in the MCA Module

Visualization Fine-Grained Alignment Fig 5 shows the alignment between visual regions and report phrases using color-coding, validating the effectiveness of hypergraph matching. For instance, the phrase "There is a screw in the right shoulder" aligns with the corresponding pink region in the image. However, some images from certain perspectives reveal incomplete correspondence between pathological phrases and their regions, highlighting the need to consider multi-view visual information to enhance alignment accuracy and comprehensiveness.

Evaluation of Zero-Shot Performance

We evaluated the zero-shot capabilities of DAMPER on the MIMIC-ABN dataset (Ni et al. 2020). Specifically, we trained the model on the MIMIC-CXR dataset and subsequently assessed its performance on the MIMIC-ABN test set. As shown in Table 6, while DAMPER performed slightly below the ICON model on BLEU-1 and ROUGE-L metrics, it demonstrated superior results on BLEU-2 through BLEU-4 and METEOR metrics. These results highlight DAMPER’s notable advantages in capturing semantic fluency and overlapping textual information. This performance underscores DAMPER’s strong generalization capability in zero-shot medical report generation tasks, providing robust evidence of its potential for clinical applications.

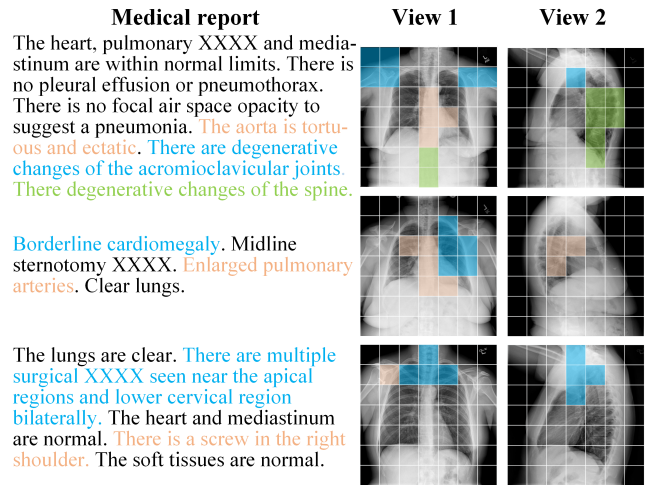


Figure 5: Visualization of alignment examples using hypergraph matching. Text phrases in the report and image regions, color-coded similarly, represent alignments established by hypergraph matching. Different colors distinguish multiple alignment relationships.

Methods	B-1	B-2	B-3	B-4	M	R-L
R2Gen (EMNLP’20)	0.290	0.157	0.093	0.061	0.105	0.208
R2GenCMN (ACL’21)	0.264	0.140	0.085	0.056	0.098	0.212
ORGAN (ACL’23)	0.314	0.180	0.114	0.078	0.120	0.234
RECAP (EMNLP’23)	0.321	0.182	0.116	0.080	0.120	0.223
ICON (EMNLP’24)	0.337	0.195	0.126	0.086	0.129	0.236
DAMPER	0.316	0.195	0.134	0.091	0.131	0.209

Table 6: Zero-shot test results of DAMPER on the MIMIC-ABN dataset.

Conclusion

This study introduces DAMPER, a dual-stage medical report generation framework inspired by the clinical report-writing process. DAMPER consists of MCG and HFG stages to simulate physicians’ overall observation and detailed examination, respectively. In MCG, the MCA module employs GAN to extract visual features corresponding to MeSH terms, which are input into the CMG to generate guiding MeSH. HFG constructs hypergraphs to model complex relationships between modality elements, designing Intra-RCA and Inter-RCA modules to generate fine-grained visual features aligning with detailed report descriptions. Finally, the outputs from MCA, CMG, and Intra-RCA are integrated into HFG’s report decoder to produce accurate medical reports. Experiments conducted on the IU-Xray and MIMIC-CXR datasets demonstrate that DAMPER outperforms state-of-the-art methods. Ablation studies validate the effectiveness of each module, while visualization results highlight DAMPER’s ability to generate reports that accurately reflect pathological conditions. Furthermore, zero-shot testing results reveal DAMPER’s strong generalization capability, as evidenced by its outstanding performance on the MIMIC-ABN test set.

Acknowledgments

This work was supported by the National Natural Science Foundation of China under Grant 82261138629 and 12326610; Guangdong Provincial Key Laboratory under Grant 2023B1212060076, and Shenzhen Municipal Science and Technology Innovation Council under Grant JCYJ20220531101412030; and in part by the project of Shenzhen Science and Technology Innovation Committee under Grant JCYJ20240813141424032, the Foundation for Young innovative talents in ordinary universities of Guangdong under Grant 2024KQNCX042.

References

- Brady, A. P. 2018. Radiology reporting—from Hemingway to HAL? *Insights into imaging*, 9: 237–246.
- Cao, Y.; Cui, L.; Yu, F.; Zhang, L.; Li, Z.; et al. 2022. KdTNet: Medical Image Report Generation via Knowledge-Driven Transformer. In *Database Systems for Advanced Applications - 27th International Conference, DASFAA*, volume 13247 of *Lecture Notes in Computer Science*, 117–132.
- Cao, Y.; Cui, L.; Zhang, L.; Yu, F.; Li, Z.; and Xu, Y. 2023. MMTN: Multi-Modal Memory Transformer Network for Image-Report Consistent Medical Report Generation. In *Thirty-Seventh AAAI Conference on Artificial Intelligence, AAAI*, 277–285.
- Chen, G.; Hou, L.; Chen, Y.; Dai, W.; Shang, L.; Jiang, X.; et al. 2023. mCLIP: Multilingual CLIP via Cross-lingual Transfer. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, 13028–13043.
- Chen, W.; Liu, J.; Liu, T.; and Yuan, Y. 2024a. Bi-VLGM: Bi-Level Class-Severity-Aware Vision-Language Graph Matching for Text Guided Medical Image Segmentation. *International Journal of Computer Vision*, 1–17.
- Chen, W.; Shen, L.; Lin, J.; Luo, J.; Li, X.; and Yuan, Y. 2024b. Fine-Grained Image-Text Alignment in Medical Imaging Enables Explainable Cyclic Image-Report Generation. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, 9494–9509.
- Chen, W.; Wang, P.; Ren, H.; Sun, L.; Li, Q.; Yuan, Y.; and Li, X. 2024c. Medical image synthesis via fine-grained image-text alignment and anatomy-pathology prompting. In *Medical Image Computing and Computer Assisted Intervention – MICCAI*, volume 15012, 240–250.
- Chen, Z.; Shen, Y.; Song, Y.; and Wan, X. 2021. Cross-modal Memory Networks for Radiology Report Generation. 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*.
- Chen, Z.; Song, Y.; Chang, T.; and Wan, X. 2020. Generating Radiology Reports via Memory-driven Transformer. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing, EMNLP*, 1439–1449.
- Demner-Fushman, D.; Kohli, M. D.; Rosenman, M. B.; Shooshan, S. E.; et al. 2016. Preparing a collection of radiology examinations for distribution and retrieval. *J. Am. Medical Informatics Assoc.*, 23(2): 304–310.
- Denkowski, M. J.; and Lavie, A. 2011. Meteor 1.3: Automatic Metric for Reliable Optimization and Evaluation of Machine Translation Systems. In *Proceedings of the Sixth Workshop on Statistical Machine Translation, WMT@EMNLP*, 85–91.
- Fu, K.; Liu, S.; Luo, X.; and Wang, M. 2021. Robust Point Cloud Registration Framework Based on Deep Graph Matching. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 8893–8902.
- Gao, Y.; Liu, J.; Xu, Z.; Wu, T.; Zhang, E.; Li, K.; Yang, J.; Liu, W.; and Sun, X. 2024. SoftCLIP: Softer Cross-Modal Alignment Makes CLIP Stronger. In *Thirty-Eighth AAAI Conference on Artificial Intelligence, AAAI*, 1860–1868.
- Gu, T.; Liu, D.; Li, Z.; and Cai, W. 2024. Complex Organ Mask Guided Radiology Report Generation. In *IEEE/CVF Winter Conference on Applications of Computer Vision (WACV)*, 7980–7989.
- Hou, W.; Cheng, Y.; Xu, K.; Hu, Y.; Li, W.; and Liu, J. 2024. ICON: Improving Inter-Report Consistency of Radiology Report Generation via Lesion-aware Mix-up Augmentation. *arXiv preprint arXiv:2402.12844*.
- Hou, W.; Cheng, Y.; Xu, K.; et al. 2023a. RECAP: Towards Precise Radiology Report Generation via Dynamic Disease Progression Reasoning. In *Findings of the Association for Computational Linguistics: EMNLP*, 2134–2147.
- Hou, W.; Xu, K.; Cheng, Y.; Li, W.; and Liu, J. 2023b. ORGAN: Observation-Guided Radiology Report Generation via Tree Reasoning. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, ACL, 8108–8122.
- Huang, X.; and Gong, H. 2024. A Dual-Attention Learning Network With Word and Sentence Embedding for Medical Visual Question Answering. *IEEE Trans. Medical Imaging*, 43(2): 832–845.
- Huang, Z.; Zhang, X.; and Zhang, S. 2023. KiUT: Knowledge-Injected U-Transformer for Radiology Report Generation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 19809–19818.
- 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*, 2607–2615.
- Jing, B.; Xie, P.; and Xing, E. 2017. On the automatic generation of medical imaging reports. *arXiv preprint arXiv:1711.08195*.
- Johnson, A. E.; Pollard, T. J.; Greenbaum, N. R.; Lungren, M. P.; et al. 2019. MIMIC-CXR-JPG, a large publicly available database of labeled chest radiographs. *arXiv preprint arXiv:1901.07042*.

- Li, Y.; Yang, B.; Cheng, X.; Zhu, Z.; Li, H.; and Zou, Y. 2023. Unify, Align and Refine: Multi-Level Semantic Alignment for Radiology Report Generation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, 2863–2874.
- Lin, C.-Y. 2004. Rouge: A package for automatic evaluation of summaries. In *Text summarization branches out*, 74–81.
- Liu, F.; Wu, X.; Ge, S.; Fan, W.; and Zou, Y. 2021. Exploring and Distilling Posterior and Prior Knowledge for Radiology Report Generation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 13753–13762.
- Liu, J.; Wang, W.; Ma, Z.; Huang, G.; SU, Y.; Chang, K.-J.; Chen, W.; Li, H.; Shen, L.; and Lyu, M. 2024a. Med-chain: Bridging the Gap Between LLM Agents and Clinical Practice through Interactive Sequential Benchmarking. *arXiv preprint arXiv:2412.01605*.
- Liu, S.; Chen, W.; Liu, J.; Luo, X.; and Shen, L. 2024b. GEM: Context-Aware Gaze Estimation with Visual Search Behavior Matching for Chest Radiograph. In *International Conference on Medical Image Computing and Computer-Assisted Intervention—MICCAI*, volume 15001, 525–535.
- Luo, X.; Liu, C.; Wong, W.; Wen, J.; Jin, X.; and Xu, Y. 2023. MVCINN: multi-view diabetic retinopathy detection using a deep cross-interaction neural network. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 37, 8993–9001.
- Luo, X.; Xu, Q.; Wang, Z.; Huang, C.; Liu, C.; Jin, X.; and Zhang, J. 2024. A Lesion-Fusion Neural Network for Multi-View Diabetic Retinopathy Grading. *IEEE J. Biomed. Health Informatics*.
- Mei, X.; Yang, L.; Gao, D.; Cai, X.; Han, J.; and Liu, T. 2024. PhraseAug: An Augmented Medical Report Generation Model with Phrasebook. *IEEE Trans. Medical Imaging*.
- Ni, J.; Hsu, C.; Gentili, A.; and McAuley, J. J. 2020. Learning Visual-Semantic Embeddings for Reporting Abnormal Findings on Chest X-rays. In *Findings of the Association for Computational Linguistics: EMNLP*, volume EMNLP 2020 of *Findings of ACL*, 1954–1960.
- Nicolson, A.; Dowling, J.; and Koopman, B. 2023. Improving chest X-ray report generation by leveraging warm starting. *Artif. Intell. Medicine*, 144: 102633.
- 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*, 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*, 448–458.
- Radford, A.; Kim, J. W.; Hallacy, C.; Ramesh, A.; Goh, G.; Agarwal, S.; et al. 2021. Learning Transferable Visual Models From Natural Language Supervision. In *Proceedings of International Conference on Machine Learning (ICML)*, volume 139, 8748–8763.
- Shin, H.-C.; Roberts, K.; Lu, L.; Demner-Fushman, D.; et al. 2016. Learning to Read Chest X-Rays: Recurrent Neural Cascade Model for Automated Image Annotation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2497–2506.
- Smit, A.; Jain, S.; Rajpurkar, P.; Pareek, A.; Ng, A. Y.; and Lungren, M. P. 2020. Combining Automatic Labelers and Expert Annotations for Accurate Radiology Report Labeling Using BERT. In *Proceedings of Conference on Empirical Methods in Natural Language Processing, EMNLP*, 1500–1519.
- Tanida, T.; Müller, P.; Kaissis, G.; and Rueckert, D. 2023. Interactive and explainable region-guided radiology report generation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 7433–7442.
- Xu, Q.; Luo, X.; Huang, C.; Liu, C.; Wen, J.; Wang, J.; and Xu, Y. 2024. HACDR-Net: Heterogeneous-Aware Convolutional Network for Diabetic Retinopathy Multi-Lesion Segmentation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, 6342–6350.
- Yan, B.; and Pei, M. 2022. Clinical-BERT: Vision-Language Pre-training for Radiograph Diagnosis and Reports Generation. In *Thirty-Sixth AAAI Conference on Artificial Intelligence, AAAI*, 2982–2990.
- 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 Analysis*, 86: 102798.
- You, D.; Liu, F.; Ge, S.; Xie, X.; Zhang, J.; and Wu, X. 2021. AlignTransformer: Hierarchical Alignment of Visual Regions and Disease Tags for Medical Report Generation. In *Medical Image Computing and Computer Assisted Intervention - MICCAI*, volume 12903, 72–82.
- Yuan, J.; Liao, H.; Luo, R.; and Luo, J. 2019. Automatic Radiology Report Generation Based on Multi-view Image Fusion and Medical Concept Enrichment. In *Medical Image Computing and Computer Assisted Intervention - MICCAI*, volume 11769, 721–729.
- Zhan, C.; Lin, Y.; Wang, G.; Wang, H.; and Wu, J. 2024. MedM2G: Unifying Medical Multi-Modal Generation via Cross-Guided Diffusion with Visual Invariant. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 11502–11512.
- Zhang, J.; Wu, S.; Liu, P.; and Shen, L. 2024. A New Dataset and Baseline Model for Rectal Cancer Risk Assessment in Endoscopic Ultrasound Videos. In *Medical Image Computing and Computer Assisted Intervention - MICCAI*, volume 15003, 564–573.
- Zhang, X.; Sun, X.; Luo, Y.; Ji, J.; Zhou, Y.; Wu, Y.; Huang, F.; and Ji, R. 2021. RSTNet: Captioning With Adaptive Attention on Visual and Non-Visual Words. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 15465–15474.
- Zhao, G.; Zhao, Z.; Gong, W.; and Li, F. 2023. Radiology report generation with medical knowledge and multilevel image-report alignment: A new method and its verification. *Artificial Intelligence in Medicine*, 146: 102714.