

FaNe: Towards Fine-Grained Cross-Modal Contrast with False-Negative Reduction and Text-Conditioned Sparse Attention

Peng Zhang¹, Zhihui Lai^{1*}, Wenting Chen^{2*}, Xu Wu¹, Heng Kong³

¹College of Computer Science and Software Engineering, Shenzhen University

²Department of Radiation Oncology, Stanford University

³Department of Thyroid and Breast Surgery, Affiliated Hospital Group of Guangdong Medical University, Shenzhen Baoan Central Hospital

zhang_cloud1@163.com, laizhihui@szu.edu.cn, wentchen@stanford.edu, csxunwu@gmail.com, generaldoc@126.com

Abstract

Medical vision-language pre-training (VLP) offers significant potential for advancing medical image understanding by leveraging paired image-report data. However, existing methods are limited by **False Negatives (FaNe)** induced by semantically similar texts and insufficient fine-grained cross-modal alignment. To address these limitations, we propose FaNe, a semantic-enhanced VLP framework. To mitigate false negatives, we introduce a semantic-aware positive pair mining strategy based on text-text similarity with adaptive normalization. Furthermore, we design a text-conditioned sparse attention pooling module to enable fine-grained image-text alignment through localized visual representations guided by textual cues. To strengthen intra-modal discrimination, we develop a hard-negative aware contrastive loss that adaptively reweights semantically similar negatives. Extensive experiments on five downstream medical imaging benchmarks demonstrate that FaNe achieves state-of-the-art performance across image classification, object detection, and semantic segmentation, validating the effectiveness of our framework.

Code — <https://github.com/Aventador8/FaNe>

1 Introduction

Representation learning for medical radiographs has gained significant attention due to the growing availability of annotated datasets (Johnson et al. 2019a; Chambon et al. 2024). Supervised methods excel in downstream tasks using large-scale labeled data (Wu et al. 2024; Chiu et al. 2025), but acquiring such annotations is labor-intensive and costly, especially in the medical domain requiring expert knowledge. To overcome this, unsupervised pre-training approaches (Zhang et al. 2023; Lai et al. 2024; Zhongyi Shui 2025) leverage free-text medical reports as auxiliary supervision, aligning images with text to capture high-level semantics without manual labels, offering a compelling alternative to supervised methods.

Although existing vision-language pre-training (VLP) methods have advanced performance on downstream medical tasks, they often overlook the false negative problem. In standard VLP training, an image paired with its corresponding report is treated as a positive sample, and all other

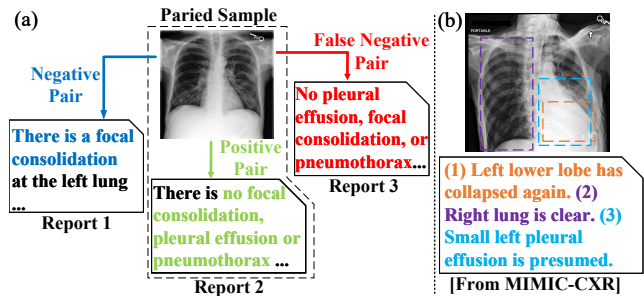


Figure 1: Illustration of false-negative problem and fine-grained alignment. (a) CLIP-style methods simply treat unpaired text or image as negatives, easily resulting in false-negative pairs (i.e, semantically similar but from different reports). (b) CLIP-style methods capture only global alignment, while color-coded sentences reveal the need for fine-grained region-level alignment.

reports are considered negative samples, as illustrated in Figure 1(a). In clinical practice, different patients may exhibit the same diseases or lesions, resulting in reports with highly similar or even identical descriptions. As shown in Figure 1(a), report 2 and report 3 contain the same findings, yet report 3 is still treated as a negative sample for the image, introducing erroneous negative pairs. In fact, the image should align with the first part of report 3, which shares similar findings with report 2, while remaining disaligned with the other part of report 3. However, current VLP methods have not considered such false negative problem, resulting in incorrect alignment during pre-training. Thus, it is essential to develop alignment strategies that can effectively mitigate the false negative issue.

Another key challenge in VLP is the limited capture of detailed visual features, as CLIP performs global image-text alignment. Current VLP approaches for fine-grained alignment typically employ cross-attention to align local image and text features or map visual and textual tokens into a shared embedding space. However, these methods often depend on coarse alignment strategies that struggle to accurately capture semantically relevant local regions within the image. As illustrated in Figure 1(b), each sentence in a report generally corresponds to specific regions within the image,

*Corresponding authors

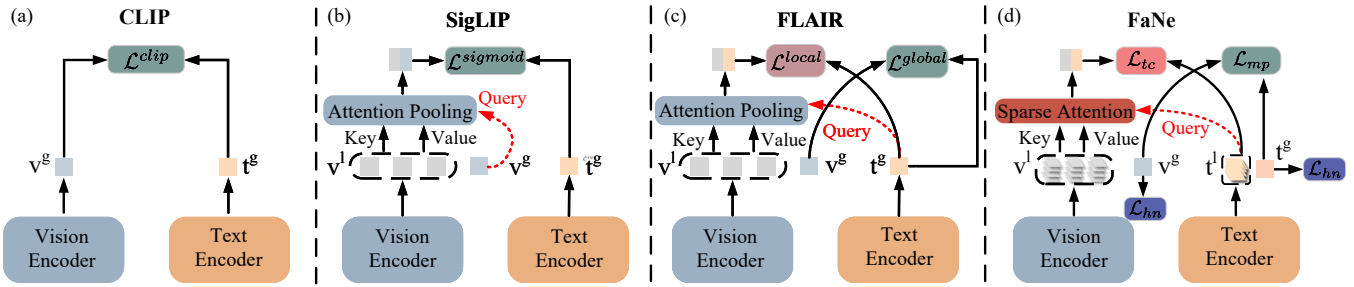


Figure 2: Comparison of text-conditioned sparse attention pooling with existing methods. (a) CLIP aligns global image v^g and text t^g tokens. (b) SigLIP uses a learnable global token v^g as query to pool local tokens v^l via cross-attention. (c) FLAIR adopts text-conditioned pooling with t^g as query to aggregate v^l for local alignment. (d) FaNe applies sparse attention for local text-aware features and employs a hard-negative aware contrastive loss \mathcal{L}_{hn} to separate hard negatives for fine-grained understanding.

underscoring the necessity for fine-grained alignment conditioned on sentence-level semantics. For example, FLAIR (Xiao et al. 2025) leverages text-conditioned attention pooling to improve alignment but struggles to accurately localize relevant image regions, as shown in Figure 2(c). This is because cross-attention alone lacks the capacity to enforce precise spatial focus. Similarly, AdaMatch (Chen et al. 2024) attempts local alignment via implicit matching of image blocks to keywords, but lacks mechanisms to guarantee alignment accuracy. Thus, it is necessary to enhance fine-grained visual understanding in VLP models, as it underpins a broad spectrum of downstream tasks, including semantic segmentation and object detection.

To address these challenges, we introduce a semantic-enhanced VLP framework incorporating four key components: semantic class division, multi-positive global alignment, text-conditioned fine-grained alignment, and hard-negative intra-modal contrast. To mitigate false negatives, our **semantic class division** module employs semantic-aware adaptive normalization to uncover latent positives and categorize negatives into hard and easy subsets. To optimize multi-positive alignment within a batch, we propose **multi-positive global alignment**, which dynamically aligns a variable number of positive pairs. For precise local alignment, we propose **text-conditioned fine-grained alignment**, leveraging text-conditioned sparse attention pooling to generate localized image embeddings directly guided by textual cues. Furthermore, to enhance intra-modal discrimination, our **hard-negative intra-modal contrast** module adaptively reweights semantically similar negative instances, improving the model’s capacity to discern subtle intra-modal distinctions. In summary, our contributions can be described as follows:

- We propose a semantic-aware positive mining strategy to mitigate false negatives by identifying latent positives via text-text similarity and adaptive normalization.
- A text-conditioned sparse attention pooling module is designed for fine-grained image-text alignment, enabling localized visual grounding.
- A hard-negative aware contrastive loss is introduced, which adaptively reweights semantically similar negatives to strengthen intra-modal discrimination.

2 Related Work

Vision-Language Pre-training. VLP leverages cross-modal interactions between vision and language to learn joint representations without explicit supervision. CLIP (Radford et al. 2021) use a contrastive loss to match global image tokens with global text tokens (Figure 2(a)). However, CLIP only learns coarse-grained representation and ignores fine-grained alignment between modalities. Many works have been proposed to build fine-grained relationships between modalities, including token-level alignment (Wang et al. 2022a), knowledge-guided multi-granularity alignment (Li et al. 2024), and soft assignments allowing many-to-many mappings (Gao et al. 2024). Along with attention pooling to form the global image embeddings (Figure 2(b)), SigLIP (Zhai et al. 2023) replaced the Softmax loss of vision-language pre-training with a Sigmoid-based loss. FLAIR (Xiao et al. 2025) employs text-conditioned attention pooling to produce fine-grained image representations (Figure 2(c)). However, each sentence in a medical report typically corresponds to only a sparse local region within the image. To capture this correspondence, we propose a text-conditioned sparse attention pooling mechanism that facilitates fine-grained visual representation learning, complemented by hard-negative aware intra-modal discrimination (Figure 2(d)).

False-Negative Problem. Simply treating unpaired images and reports as negative pairs can lead to issues of false-negative and degrade the quality of representation. MedCLIP (Wang et al. 2022b) replace the InfoNCE loss with semantic matching loss based on medical knowledge to eliminate false negatives. MLIP (Li et al. 2024) employs a knowledge-guided class-level contrastive learning strategy in conjunction with a divergence encoder to distinguish false negatives from true negatives. SAT (Liu et al. 2023a) computes semantic similarity between reports to group image-report pairs into positive, negative, and neutral sets, thereby avoiding the misclassification of semantically similar but unpaired samples as negatives. However, the fixed normalization of semantic similarity computation in SAT is suboptimal for varying mini-batches and different training stages. To alleviate this problem, we propose a semantic-enhanced adaptive normalization strategy to eliminate struc-

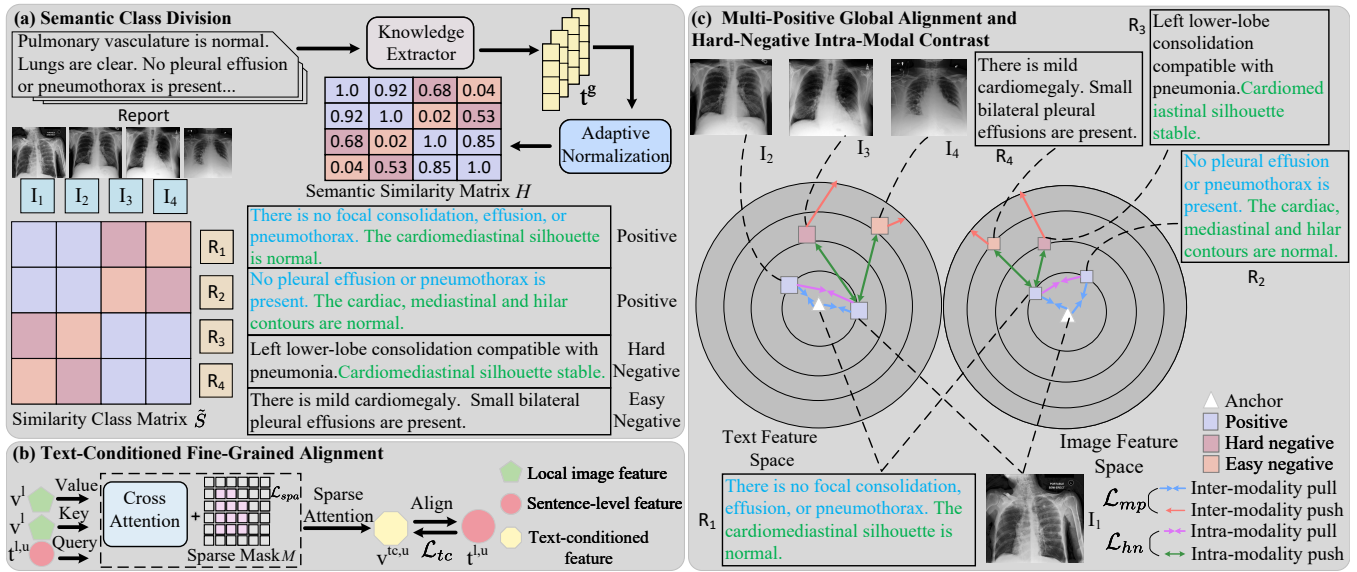


Figure 3: Overview of FaNe. (a) A semantic similarity matrix is generated via a knowledge encoder (e.g., BioClinicalBERT (Alsentzer et al. 2019)), enabling image-text pairs to be classified by a similarity threshold, with negatives further divided into hard and easy cases. (b) Text-conditioned features are derived through sparse attention, enabling alignment with sentence-level representations. (c) The combined losses \mathcal{L}_{mp} and \mathcal{L}_{hn} shape the feature space, where longer arrows indicate stronger repulsion for fine-grained concept learning. Sentences with the same highlight color share semantic meaning.

tured common semantics and stabilize training.

3 Method

Our overall workflow is illustrated in Figure 3. Given reports t , we first use a pretrained knowledge extractor f_{tool} to encode them into global text representations t^g . Next, we compute the semantic similarity matrix \tilde{S} via the semantic-enhanced adaptive normalization, followed by constructing the similarity class matrix H using a predefined threshold to identify positive and negative samples. The global alignment is then optimized by applying the loss \mathcal{L}_{mp} . We further extract local text representations $t^{l,u}$, which are fed into a cross-attention mechanism to obtain text-conditioned visual representations $v^{tc,u}$. The fine-grained alignment loss \mathcal{L}_{tc} and the sparse regularization loss \mathcal{L}_{spa} are jointly applied to align t^l with $v^{tc,u}$. Finally, for intra-modal discrimination, the hard-negative aware intra-modal contrastive loss \mathcal{L}_{hn} is applied to the global image representations v^g and global text representations t^g , enhancing the model’s ability to distinguish subtle semantic differences within each modality.

3.1 Semantic Class Division

Given a training set $D = \{(v_1, t_1), \dots, (v_N, t_N)\}$, where v_i denotes a radiograph and t_i is its corresponding report, the objective is to not only leverage ground-truth pairs (v_i, t_i) but also to correct false negatives (v_i, t_j) arising from semantically similar samples. A pretrained knowledge extractor f_{tool} (e.g., BioClinicalBERT (Alsentzer et al. 2019)) encodes each report into global and local representations: $f_{tool}(t_i) = [t_i^l, t_i^g]$. To reduce bias from semantic redun-

dancy in clinical narratives and stabilize similarity computation across batches, we introduce a semantic-enhanced adaptive normalization scheme that calibrates report-to-report similarity via batch statistics and exponential moving average (EMA) smoothing:

$$p_b = \frac{1}{B} \sum_{i=1}^B \tilde{t}_i^g \quad (1)$$

where B is the batchsize and \tilde{t}_i^g denotes unit vectors $\frac{t_i^g}{\|t_i^g\|_2}$. The batch’s semantic base similarity is measured as the mean cosine similarity between each report and prototype:

$$\hat{o}_t^* = \frac{1}{B} \sum_{i=1}^B \langle \tilde{t}_i^g, p_b \rangle \quad (2)$$

where $\langle \cdot, \cdot \rangle$ represents cosine similarity function. To improve training stability and prevent sudden shifts across batches, we smooth the estimate using EMA:

$$o_t^* = \alpha \cdot \hat{o}_t^* + (1 - \alpha) \cdot \hat{o}_{t-1}^* \quad (3)$$

where α is set to 0.05, and $o_0^* = \hat{o}_0^*$. Given the normalized pairwise cosine similarity matrix $S = \langle \tilde{t}_i^g, \tilde{t}_j^g \rangle$, we perform center-shift normalization to obtain semantic similarity matrix \tilde{S} :

$$\tilde{S} = \frac{S - o_t^*}{1 - o_t^* + \epsilon} \quad (4)$$

where ϵ is a small constant to avoid division by zero.

After obtaining \tilde{S} , we define the similarity class matrix $H \in \{0, 1\}^{B \times B}$ as follows:

$$H = \mathbb{1} [\tilde{S} > \kappa] \quad (5)$$

where κ is a predefined threshold. Note that H re-assigns a variable number of positives to each sample.

3.2 Multi-Positive Global Alignment

The symmetrical contrastive loss employed in CLIP, which facilitates bidirectional text-to-image and image-to-text alignment, is designed to handle only a single positive pair per sample. This design choice conflicts with the requirements outlined in Section 3.1, where the proposed methods necessitate training with a variable number of positive pairs per sample. To address this limitation, we adopt a sigmoid-based contrastive loss following SigLIP (Zhai et al. 2023), which naturally accommodates multiple positive pairs within a batch. Compared to the InfoNCE loss (Oord, Li, and Vinyals 2018), this formulation offers greater flexibility for handling many-to-many alignments. Hence, multi-positive global alignment loss is defined:

$$\mathcal{L}_{mp} = -\frac{1}{N} \sum_{i=1}^N \sum_{j=1}^N \log \frac{1}{1 + e^{h_{ij}(-\langle v_i^g, v_j^g \rangle / \tau_1 + b)}} \quad (6)$$

where h_{ij} denote the (i, j) -th entry of pairwise label matrix H (-1 for negative and 1 for positive pairs). To modulate similarity scaling, a temperature τ_1 is introduced. Additionally, a learnable bias term b , initialized to a negative value, is incorporated to offset the imbalance between positive and negative pairs, ensuring a lower initial loss by reducing the contribution of the overwhelming number of negatives during early training.

3.3 Text-Conditioned Fine-Grained Alignment

A key distinction between medical and natural images lies in the critical role of multi-scale visual features. Medical diagnosis relies on both global anatomical context and fine-grained local cues, demanding effective integration of hierarchical representations. Moreover, medical reports often comprise multiple sentences, each describing region-specific findings (as shown in Figure 1(b)), further underscoring the need for localized visual-textual alignment. Therefore, we propose to contextualize the image representations with the individual sentences, producing a unique image representation for every image-text pair.

Image and Text Representation. Two independent encoders f_{img} and f_{txt} are used to extract features:

$$f_{img}(v) = [v^l, v^g], f_{txt}(t) = [t^l, t^g] \quad (7)$$

where $v^l \in \mathbb{R}^{I \times D}$ and $t^l \in \mathbb{R}^{P \times L \times D}$ represent local image embeddings and text embeddings, respectively. $v^g \in \mathbb{R}^D$ and $t^g \in \mathbb{R}^D$ denote global image embeddings and text embeddings, respectively. I is the number of local image patches. P is the total number of sentences in the report. L is the maximum text length and D refers to the dimension of features. In our experiments, the dimension of final output from both images and texts are the same.

Sparse Attention Mask. Each sentence in a medical report typically corresponds to only a local region of the image. Hence, how to effectively model fine-grained semantic

alignment is a unique challenge. To address it, a learnable sparse attention mask is devised:

$$M = \sigma(MLP([\mathcal{R}(v^l), \mathcal{R}(t^{l,u})])) \quad (8)$$

where u is the sentence index of each report and $M \in \mathbb{R}^{L \times I}$. $MLP(\cdot)$ denotes two layer linear transformations. \mathcal{R} represents reshape operation. σ is the sigmoid activation function.

To promote sparsity and reduce redundancy, we further impose the sparsity constraint overlaid on top of M :

$$\mathcal{L}_{spa} = \sum_{i=1}^L \sum_{j=1}^I |M_{i,j}| \quad (9)$$

where $M_{i,j}$ are the activation values of the learnable attention mask M .

Cross-Attention Fine-Grained Alignment. To infuse semantic context into visual representations, we employ a text-conditioned cross-attention sparse pooling mechanism. Specifically, the local sentence embedding serves as a query to aggregate local image patch embeddings, yielding a text-conditioned visual representation $v^{tc,u}$ as formulated:

$$v^{tc,u} = \sum_{k=1}^I LN\left(\left(\sigma\left(\frac{t^{l,u} W_q (v^{l,k} W_k)^T}{\sqrt{D}}\right) \cdot M\right) v^{l,k} W_v W_o\right) \quad (10)$$

where W_q, W_k, W_v, W_o are learnable matrices. k is the spatial location index of each patch. LN stands for LayerNorm.

After obtaining the $v^{tc,u}$, the InfoNCE loss is applied on $(v^{tc,u}, t^{l,u})$ pairs where $v^{tc,u}$ and $t^{l,u}$ from the same sentence are regarded as a positive pair, while those from different sentences are negative pairs. Given that semantically similar sentences frequently occur across different reports, we restrict the selection of negative samples for local alignment to sentences within the same report, thereby mitigating false negatives and enhancing alignment precision. We symmetrize the optimization by jointly minimizing the local text-to-image and image-to-text contrastive losses, with the total loss defined as the sum of both directional components:

$$\mathcal{L}_{tc} = (\mathcal{L}_{t2i} + \mathcal{L}_{i2t}) / 2 \quad (11)$$

$$\mathcal{L}_{t2i} = -\frac{1}{\sum_{i=1}^N P_i} \sum_{i=1}^N \sum_{u=1}^{P_i} \log \frac{\exp(t_i^{l,u} \cdot (v_i^{tc,u})^T / \tau_2)}{\sum_{k=1}^{P_i} \exp(t_i^{l,k} \cdot (v_i^{tc,k})^T / \tau_2)} \quad (12)$$

$$\mathcal{L}_{i2t} = -\frac{1}{\sum_{i=1}^N P_i} \sum_{i=1}^N \sum_{u=1}^{P_i} \log \frac{\exp(t_i^{l,u} \cdot (v_i^{tc,u})^T / \tau_2)}{\sum_{k=1}^{P_i} \exp(t_i^{l,k} \cdot (v_i^{tc,k})^T / \tau_2)} \quad (13)$$

where τ_2 is the temperature parameter and P_i denotes the number of sentences of i -th report.

3.4 Hard-Negative Intra-Modal Contrast

The core motivation for introducing the hard-negative intra-modal contrastive loss lies in its ability to explicitly strengthen the model’s discriminative capacity for subtle semantic variations. By assigning greater importance to semantically similar negative samples, which are often under-emphasized in standard contrastive learning, this loss formulation encourages the model to focus on fine-grained intra-modal distinctions. Such targeted supervision is particularly

crucial in medical vision language pretraining, where accurately capturing nuanced clinical differences can significantly impact downstream diagnostic performance. So the symmetric hard-negative intra-modal contrastive loss is defined as follow:

$$\mathcal{L}_{hn} = (\mathcal{L}_{img} + \mathcal{L}_{text})/2 \quad (14)$$

$$\mathcal{L}_{img} = -\frac{1}{N} \sum_{i=1}^N \log \frac{1}{\sum_{j=1}^N \exp(y_{ij} \cdot \alpha_{ij} \cdot (\mathbf{v}_i^g \cdot (\mathbf{v}_j^g)^T / \tau_3))} \quad (15)$$

$$\mathcal{L}_{text} = -\frac{1}{N} \sum_{i=1}^N \log \frac{1}{\sum_{j=1}^N \exp(y_{ij} \cdot \beta_{ij} \cdot (\mathbf{t}_i^g \cdot (\mathbf{t}_j^g)^T / \tau_3))} \quad (16)$$

$$y_{ij} = \begin{cases} 1 & \text{if } h_{ij} = -1 \\ 0 & \text{if } h_{ij} = 1 \end{cases} \quad (17)$$

$$\alpha_{ij} = \frac{y_{ij} \cdot \mathbf{v}_i^g \cdot (\mathbf{v}_j^g)^T / \tau_3}{\sum_{k \neq i} y_{ik} \cdot \mathbf{v}_i^g \cdot (\mathbf{v}_k^g)^T / \tau_3}, \beta_{ij} = \frac{y_{ij} \cdot \mathbf{t}_i^g \cdot (\mathbf{t}_j^g)^T / \tau_3}{\sum_{k \neq i} y_{ik} \cdot \mathbf{t}_i^g \cdot (\mathbf{t}_k^g)^T / \tau_3} \quad (18)$$

where y_{ij} stands for the index of negative pairs and τ_3 is the temperature parameter. Weights α_{ij} and β_{ij} are designed that hard-negative pairs (i.e., higher semantic similarity in negative pairs) are underscored. This enables the model to effectively capture fine-grained semantic distinctions between true positive pairs and challenging hard-negative counterparts, thereby fostering more discriminative and concept-aware representations.

3.5 Overall Objective

The overall loss function of the proposed FaNe contains the following terms: multi-positive global alignment loss \mathcal{L}_{mp} , sparse regularization loss \mathcal{L}_{spa} , text-conditioned fine-grained alignment loss \mathcal{L}_{tc} and hard-negative aware intra-modal contrastive loss \mathcal{L}_{hn} . We optimize these terms jointly via a weighted sum:

$$\mathcal{L} = \mathcal{L}_{mp} + \lambda_1 \mathcal{L}_{tc} + \lambda_2 \mathcal{L}_{hn} + \lambda_3 \mathcal{L}_{spa} \quad (19)$$

where λ_1 , λ_2 , and λ_3 serve as hyper-parameters that modulate the relative contribution of each loss component.

4 Experiments

We pre-train the proposed FaNe framework on a large-scale medical image–report corpus, and comprehensively assess the quality of the learned visual representations across five benchmark datasets spanning three critical downstream medical imaging tasks. Furthermore, extensive ablation studies are conducted to systematically evaluate the contribution of each core component.

4.1 Pre-Training Setup

Dataset. We pre-train the proposed FaNe framework on the MIMIC-CXR v2 dataset (Johnson et al. 2019b), which comprises 377,110 chest radiographs and 227,827 corresponding medical reports. Following the preprocessing protocol in PRIOR (Cheng et al. 2023), we retain only frontal-view images to align with the evaluation setup used in downstream

tasks, discarding all lateral-view images. Furthermore, to ensure sufficient textual content for semantic supervision, we filter out reports with fewer than four sentences. After pre-processing, we obtain a curated dataset of 182,475 high-quality image–report pairs for pre-training.

Implementation Details. Following previous works, we use BioClinicalBERT (Alsentzer et al. 2019) as the text encoder. For image feature encoding, we apply ResNet50 (He et al. 2016) as the backbone. We resize all images into 224×224 , and train our framework 50 epochs on 2 pieces of RTX 4090 GPUs with batch size of 98. The dimension D is set 128 and the temperature hyper-parameters are $\tau_1 = 0.1$, $\tau_2 = 0.07$, $\tau_3 = 0.07$. The learning rate is initialized as $4e-4$ and decays following a cosine policy. We set $\lambda_1 = 1$, $\lambda_2 = 1$, $\lambda_3 = 1$. More details can be found in the supplementary material.

4.2 Downstream Tasks

Medical Semantic Segmentation. We evaluate the segmentation performance of our model on the **RSNA** Pneumonia (Shih et al. 2019) and **SIIM** Pneumothorax (Zawacki et al. 2019) datasets. Following the GLoRIA (Huang et al. 2021), we fine-tune a U-Net (Ronneberger, Fischer, and Brox 2015) with a frozen ResNet-50 encoder initialized from our pre-trained model, while training the decoder on varying fractions of labeled data (1%, 10%, and 100%). Segmentation quality is measured using the Dice similarity coefficient.

Medical Image Classification. We evaluate the classification performance of our model on three standard benchmarks: **RSNA** Pneumonia (Shih et al. 2019), **COVIDx** (Wang, Lin, and Wong 2020), and **CheXpert** (Irvin et al. 2019). Following the linear evaluation protocol from MGCA (Wang et al. 2022a), we freeze the pre-trained ViT-B/16 (Dosovitskiy et al. 2020) or ResNet-50 backbone and train a linear classifier for downstream prediction. To assess data efficiency, experiments are conducted using 1%, 10%, and 100% of the training data. Performance is measured by AU-ROC for RSNA and CheXpert, and classification accuracy for COVIDx, consistent with ConVIRT (Zhang et al. 2022). Further details are included in the supplementary material.

Medical Object Detection. We evaluate object detection capabilities of our pre-trained image encoder on the **RSNA** Pneumonia dataset (Shih et al. 2019) and the **Object-CXR** dataset (Healthcare 2020). Following the YOLOv3 (Redmon and Farhadi 2018) frozen backbone protocol, our pre-trained ResNet-50 encoder is fixed and integrated as the feature extractor, while only the detector head is fine-tuned. To investigate data efficiency, we conduct experiments using 1%, 10%, and 100% of the training data. Detection performance is measured using mean Average Precision (mAP), averaged over IoU thresholds ranging from 0.4 to 0.75.

4.3 Results

The results of competing methods are cited from their respective original publications. For a fair comparison, we follow the MGCA protocol to pre-train GLoRIA on the MIMIC-CXR dataset, as indicated by the \S symbol.

Results on Medical Semantic Segmentation. We evaluate the effect of fine-grained local representations on segmentation using the ResNet-50 U-Net across the RSNA

Dataset Method	FG	Semantic Segmentation (Dice)						Object Detection (mAP)					
		RSNA			SIIM			RSNA			Object CXR		
		1%	10%	100%	1%	10%	100%	1%	10%	100%	1%	10%	100%
Random Init		6.9	10.6	18.5	9.0	28.6	54.3	1.0	4.0	8.9	~	0.5	4.4
ImageNet Init		34.8	39.9	64.0	10.2	35.5	63.5	3.6	8.0	15.7	~	2.9	8.3
ConVIRT(Zhang et al. 2022)		55.0	67.4	67.5	25.0	43.2	59.9	8.2	15.6	17.9	~	8.6	15.9
GLoRIA [§] (Huang et al. 2021)	✓	60.3	68.7	68.3	37.4	57.1	64.0	11.6	16.1	24.8	~	8.9	16.6
MGCA (Wang et al. 2022a)	✓	63.0	68.3	69.8	49.7	59.3	64.2	12.9	16.8	24.9	~	12.1	19.2
M-FLAG (Liu et al. 2023b)		64.6	69.7	70.5	52.5	61.2	64.8	13.7	17.5	25.4	~	12.4	19.3
MedKLIP* (Wu et al. 2023)		66.2	69.4	71.9	50.2	60.8	63.9	8.9	16.3	24.5	~	7.1	11.6
PRIOR (Cheng et al. 2023)	✓	66.4	68.3	72.7	51.2	59.7	66.3	15.6	18.5	25.2	2.9	15.2	19.8
MLIP (Li et al. 2024)	✓	67.7	68.8	73.5	51.6	60.8	68.1	17.2	19.1	25.8	4.6	17.4	20.2
IMITATE (Liu et al. 2024)		70.5	71.4	73.8	53.9	61.7	64.5	15.3	19.7	26.4	3.9	12.7	20.3
FaNe (ours)	✓	69.5	72.4	74.1	54.1	62.3	68.8	16.4	20.6	27.2	4.9	15.6	21.2

Table 1: Results of semantic segmentation on RSNA and SIIM datasets and object detection on RSNA and Object CXR datasets. Method marked with * uses extra annotated data and the "∼" denotes mAP is smaller than 1%. The "FG" means these methods have or do not have fine-grained alignment. The § symbol denotes models pretrained on the MIMIC-CXR dataset.




Figure 4: Representative cross-modality attention maps. (a) The related sentence is "Mild thoracic scoliosis noted." (b) The related sentence is "Heart size is normal."

and SIIM datasets, as shown in Table 1. While prior methods (e.g., GLoRIA, MGCA, PRIOR, MedKLIP, MLIP) incorporate localized features during pre-training, our FaNe framework consistently outperforms them. This improvement stems from the proposed text-conditioned sparse attention pooling, which enables more precise region-level feature extraction essential for accurate segmentation.

Results on Medical Object Detection. To further validate the effectiveness of our proposed FaNe framework, we perform object detection experiments utilizing a ResNet-50-YOLOv3 architecture on the RSNA and Object CXR datasets. As presented in Table 1, FaNe consistently surpasses competing methods across most scenarios. We hypothesize that this superior performance stems from the strategic mining of potential positive samples through semantically aligned textual cues, which mitigates the adverse effects of misclassified negative samples. Consequently, this approach yields more precise and robust visual feature representations, enhancing the reliability of bounding box regression for object detection tasks in medical imaging.

Results on Medical Image Classification. We perform supervised classification on the RSNA, CheXpert, and COVIDx datasets. Specifically, we pre-train our proposed FaNe model using both ResNet-50 and ViT-B/16 architectures to evaluate the impact of different image encoders. As reported in Table 2, FaNe with the ViT-B/16 encoder

achieves superior performance across multiple datasets, suggesting that Transformer-based architectures may be particularly effective for aligning vision and language modalities in medical imaging tasks. Notably, while existing approaches achieve promising performance, they overlook hard-negative mining, leading to suboptimal representation learning. In contrast, our proposed hard-negative aware contrastive loss imposes stronger separation on semantically similar samples, encouraging the model to learn clearer semantic boundaries. This helps reduce misclassification of boundary-ambiguous categories and enhances both robustness and recognition capability.

4.4 Ablation Study

Visualization. To demonstrate the effectiveness and interpretability of our text-conditioned fine-grained alignment, we visualize cross-modality attention maps generated from image-text pairs. Using Equation 10, we compute the attention maps and resize them for overlay on the input images. As shown in Figure 4, our sparse attention pooling module with ViT as the encoder accurately highlights semantically relevant regions such as thoracic scoliosis and the heart location based on the corresponding sentence. Additional visualizations using ResNet are presented in the appendix.

Effectiveness of Sparse Attention Mask. To assess the impact of incorporating a sparse attention mask into the cross-attention mechanism, we conduct an ablation study presented in Table 3. The first row reports the baseline model without any learnable attention mask. The second row introduces a learnable attention mask without applying sparsity regularization. The final row corresponds to our full method with both learnable attention and a sparsity constraint. The results clearly demonstrate that both components contribute significantly to performance gains in semantic segmentation and object detection, validating the effectiveness of our fine-grained alignment strategy.

Effectiveness of Semantic-Enhanced Adaptive Normalization. We conduct an ablation study by toggling this com-

Dataset Method	FG	CheXpert (AUC)			RSNA (AUC)			COVIDx (ACC)		
		1%	10%	100%	1%	10%	100%	1%	10%	100%
Random Init		56.1	62.6	65.7	58.9	69.4	74.1	50.5	60.3	70.0
ImageNet Init		74.4	79.7	81.4	74.9	74.5	76.3	64.8	78.8	86.3
ConVIRT (Zhang et al. 2022)		85.9	86.8	87.3	77.4	80.1	81.3	72.5	82.5	92.0
GLoRIA [§] (Huang et al. 2021)	✓	87.1	88.7	88.0	87.0	89.4	90.2	66.5	80.5	88.8
MGCA (Wang et al. 2022a)	✓	87.6	88.0	88.2	88.6	89.1	89.9	72.0	83.5	90.5
PRIOR (Cheng et al. 2023)	✓	87.6	88.6	88.8	88.9	89.5	90.5	72.3	84.7	91.0
MedKLIP* (Wu et al. 2023)		86.2	86.5	87.7	87.3	88.0	89.3	74.5	85.2	90.3
MLIP (Li et al. 2024)	✓	87.8	88.7	88.9	88.8	89.6	90.6	73.0	85.0	90.8
MedUnifer (Zhang et al. 2025)		-	-	-	87.6	88.8	91.7	76.8	88.3	93.5
SENSE (Liu, Lu, and Wang 2024)		87.6	88.5	89.3	88.0	89.6	90.7	-	-	-
FaNe (Ours, ResNet-50)	✓	88.2	89.1	89.9	88.9	89.8	92.6	78.2	89.1	94.0
FaNe (Ours, ViT-B/16)	✓	89.7	90.4	90.8	89.3	90.2	93.1	79.5	90.7	95.5

Table 2: Image classification performance with fine-tuning on 1%, 10%, and 100% of training data for CheXpert, RSNA, and COVIDx datasets. Method marked with * uses extra annotated data. The "FG" means these methods have or do not have fine-grained alignment. The § symbol denotes models pretrained on the MIMIC-CXR dataset.

Learnable Mask	\mathcal{L}_{spa}	RSNA (Dice)			Object CXR (mAP)		
		1%	10%	100%	1%	10%	100%
		66.6	68.9	71.2	2.2	12.6	18.5
✓		67.7	70.3	72.6	3.4	13.8	19.9
✓	✓	69.5	72.4	74.1	4.9	15.6	21.2

Table 3: Ablation study on learnable sparse attention mask and sparse regularization loss.

Semantic Normalization	RSNA (Dice)			Object CXR (mAP)		
	1%	10%	100%	1%	10%	100%
	67.1	69.7	71.5	2.5	12.9	18.7
✓	69.5	72.4	74.1	4.9	15.6	21.2

Table 4: Ablation study on semantic-enhanced adaptive normalization.

ponent on and off. As shown in Table 4, semantic normalization consistently improves performance across both segmentation and detection tasks under different training data ratios. These consistent gains indicate that normalization mechanism effectively mitigates the influence of semantically similar false negatives, enhancing representation learning.

Effectiveness of Losses. As shown in Table 5, training with the semantic-aware multi-positive global alignment loss (\mathcal{L}_{mp}) alone provides a strong baseline by alleviating false negatives caused by semantically similar reports. Adding the hard-negative aware contrastive loss (\mathcal{L}_{hn}) further improves performance, demonstrating its role in enhancing intra-modal discrimination by pushing apart semantically close but incorrect negatives. Incorporating the text-conditioned sparse contrastive loss and its sparsity regularization ($\mathcal{L}_{tc} + \mathcal{L}_{spa}$) leads to additional gains, particularly in low data ratio, validating its ability to enable fine-grained and text-guided local alignment. The combination of all components yields the highest Dice scores across both datasets, confirming the complementary benefits of our design in addressing false negatives, improving alignment precision, and refining intra-modal semantic separation.

Components			RSNA (Dice)			SIIM (Dice)		
\mathcal{L}_{mp}	\mathcal{L}_{hn}	$\mathcal{L}_{tc} + \mathcal{L}_{spa}$	1%	10%	100%	1%	10%	100%
✓			66.7	69.6	71.0	50.9	59.3	64.9
✓	✓		67.8	70.8	72.3	52.1	60.5	66.1
✓		✓	68.2	71.3	73.1	52.9	61.2	67.6
✓	✓	✓	69.5	72.4	74.1	54.1	62.3	68.8

Table 5: Ablation study on losses. \mathcal{L}_{tc} and \mathcal{L}_{spa} are fixed combination as reported in Table 3.

RSNA (Dice)	Threshold				
	0.6	0.7	0.8	0.9	0.95
1%	65.7	66.6	67.7	69.1	69.5
10%	68.4	69.3	70.8	71.8	72.4
100%	70.5	70.9	71.2	72.2	74.1

Table 6: Ablation study on semantic threshold κ .

Effectiveness of Threshold. As shown in Table 6, model performance consistently improves as the semantic threshold κ increases from 0.6 to 0.95, indicating that a higher threshold helps mitigate the impact of false-negative pairs and enhances the model’s discriminative capability. The best segmentation results are achieved when κ is set to 0.95.

5 Conclusion

We present a semantic-enhanced framework to address key challenges in medical VLP, including false negatives from semantically similar reports, coarse image-text alignment, and insufficient intra-modal discrimination. Our approach integrates a semantic-aware positive pair mining strategy, a text-conditioned sparse attention pooling module, and a hard-negative aware contrastive loss to enhance the quality and granularity of learned representations. Extensive experiments on five downstream medical imaging tasks, covering classification, detection, and segmentation, demonstrate the superiority of our method over existing VLP baselines.

Acknowledgements

This work was supported in part by the Natural Science Foundation of China under Grant 62476175 and Grant 62272319, and in part by the Natural Science Foundation of Guangdong Province (Grant 2024A1515011637, 2023B1212060076), the Science and Technology Planning Project of Shenzhen Municipality under Grants JCYJ20220818095803007 and JCYJ20250604145234045, and the Special Project for Clinical and Basic Sci&Tech Innovation of Guangdong Medical University under Grant GDMULCJC2024151.

References

- Alsentzer, E.; Murphy, J.; Boag, W.; Weng, W.-H.; Jindi, D.; Naumann, T.; and McDermott, M. 2019. Publicly Available Clinical BERT Embeddings. In *Proceedings of the 2nd Clinical Natural Language Processing Workshop*, 72–78.
- Chambon, P.; Delbrouck, J.-B.; Sounack, T.; Huang, S.-C.; Chen, Z.; Varma, M.; Truong, S. Q.; Langlotz, C. P.; et al. 2024. Chexpert plus: Hundreds of thousands of aligned radiology texts, images and patients. *arXiv e-prints*, arXiv:2405.
- Chen, W.; Shen, L.; Lin, J.; Luo, J.; Li, X.; and Yuan, Y. 2024. 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.
- Cheng, P.; Lin, L.; Lyu, J.; Huang, Y.; Luo, W.; and Tang, X. 2023. Prior: Prototype representation joint learning from medical images and reports. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 21361–21371.
- Chiu, T.-M.; Li, Y.-C.; Chi, I.-C.; and Tseng, M.-H. 2025. AI-Driven Enhancement of Skin Cancer Diagnosis: A Two-Stage Voting Ensemble Approach Using Dermoscopic Data. *Cancers*, 17(1): 137.
- Dosovitskiy, A.; Beyer, L.; Kolesnikov, A.; Weissenborn, D.; Zhai, X.; Unterthiner, T.; Dehghani, M.; Minderer, M.; Heigold, G.; Gelly, S.; et al. 2020. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929*.
- 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 *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, 1860–1868.
- He, K.; Zhang, X.; Ren, S.; and Sun, J. 2016. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 770–778.
- Healthcare, J. 2020. Object-cxr-automatic detection of foreign objects on chest x-rays. *MIDL*.
- Huang, S.-C.; Shen, L.; Lungren, M. P.; and Yeung, S. 2021. Gloria: A multimodal global-local representation learning framework for label-efficient medical image recognition. In *Proceedings of the IEEE/CVF international conference on computer vision*, 3942–3951.
- Irvin, J.; Rajpurkar, P.; Ko, M.; Yu, Y.; Ciurea-Illcus, S.; Chute, C.; Marklund, H.; Haghgoo, B.; Ball, R.; Shpankaya, K.; et al. 2019. Chexpert: A large chest radiograph dataset with uncertainty labels and expert comparison. In *Proceedings of the AAAI conference on artificial intelligence*, volume 33, 590–597.
- Johnson, A. E.; Pollard, T. J.; Berkowitz, S. J.; Greenbaum, N. R.; Lungren, M. P.; Deng, C.-y.; Mark, R. G.; and Horng, S. 2019a. MIMIC-CXR, a de-identified publicly available database of chest radiographs with free-text reports. *Scientific data*, 6(1): 317.
- Johnson, A. E.; Pollard, T. J.; Greenbaum, N. R.; Lungren, M. P.; Deng, C.-y.; Peng, Y.; Lu, Z.; Mark, R. G.; Berkowitz, S. J.; and Horng, S. 2019b. MIMIC-CXR-JPG, a large publicly available database of labeled chest radiographs. *arXiv preprint arXiv:1901.07042*.
- Lai, H.; Yao, Q.; Jiang, Z.; Wang, R.; He, Z.; Tao, X.; and Zhou, S. K. 2024. Carzero: Cross-attention alignment for radiology zero-shot classification. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 11137–11146.
- Li, Z.; Yang, L. T.; Ren, B.; Nie, X.; Gao, Z.; Tan, C.; and Li, S. Z. 2024. Mlip: Enhancing medical visual representation with divergence encoder and knowledge-guided contrastive learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 11704–11714.
- Liu, B.; Lu, D.; Wei, D.; Wu, X.; Wang, Y.; Zhang, Y.; and Zheng, Y. 2023a. Improving medical vision-language contrastive pretraining with semantics-aware triage. *IEEE Transactions on Medical Imaging*, 42(12): 3579–3589.
- Liu, B.; Lu, Z.; and Wang, Y. 2024. Towards Medical Vision-Language Contrastive Pre-training via Study-Oriented Semantic Exploration. In *Proceedings of the 32nd ACM International Conference on Multimedia*, 4861–4870.
- Liu, C.; Cheng, S.; Chen, C.; Qiao, M.; Zhang, W.; Shah, A.; Bai, W.; and Arcucci, R. 2023b. M-flag: Medical vision-language pre-training with frozen language models and latent space geometry optimization. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, 637–647. Springer.
- Liu, C.; Cheng, S.; Shi, M.; Shah, A.; Bai, W.; and Arcucci, R. 2024. Imitate: Clinical prior guided hierarchical vision-language pre-training. *IEEE Transactions on Medical Imaging*.
- Oord, A. v. d.; Li, Y.; and Vinyals, O. 2018. Representation learning with contrastive predictive coding. *arXiv preprint arXiv:1807.03748*.
- Radford, A.; Kim, J. W.; Hallacy, C.; Ramesh, A.; Goh, G.; Agarwal, S.; Sastry, G.; Askell, A.; Mishkin, P.; Clark, J.; et al. 2021. Learning transferable visual models from natural language supervision. In *International conference on machine learning*, 8748–8763. PmLR.
- Redmon, J.; and Farhadi, A. 2018. Yolov3: An incremental improvement. *arXiv preprint arXiv:1804.02767*.

- Ronneberger, O.; Fischer, P.; and Brox, T. 2015. U-net: Convolutional networks for biomedical image segmentation. In *Medical image computing and computer-assisted intervention—MICCAI 2015: 18th international conference, Munich, Germany, October 5-9, 2015, proceedings, part III 18*, 234–241. Springer.
- Shih, G.; Wu, C. C.; Halabi, S. S.; Kohli, M. D.; Prevedello, L. M.; Cook, T. S.; Sharma, A.; Amorosa, J. K.; Arteaga, V.; Galperin-Aizenberg, M.; et al. 2019. Augmenting the national institutes of health chest radiograph dataset with expert annotations of possible pneumonia. *Radiology: Artificial Intelligence*, 1(1): e180041.
- Wang, F.; Zhou, Y.; Wang, S.; Vardhanabhuti, V.; and Yu, L. 2022a. Multi-granularity cross-modal alignment for generalized medical visual representation learning. In *Proceedings of International Conference on Neural Information Processing Systems*, 33536–33549.
- Wang, L.; Lin, Z. Q.; and Wong, A. 2020. Covid-net: A tailored deep convolutional neural network design for detection of covid-19 cases from chest x-ray images. *Scientific reports*, 10(1): 19549.
- Wang, Z.; Wu, Z.; Agarwal, D.; and Sun, J. 2022b. Med-clip: Contrastive learning from unpaired medical images and text. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing. Conference on Empirical Methods in Natural Language Processing*, volume 2022, 3876.
- Wu, C.; Zhang, X.; Zhang, Y.; Wang, Y.; and Xie, W. 2023. Medklip: Medical knowledge enhanced language-image pre-training for x-ray diagnosis. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 21372–21383.
- Wu, X.; Lai, Z.; Zhou, J.; Hou, X.; Pedrycz, W.; and Shen, L. 2024. Light-aware contrastive learning for low-light image enhancement. *ACM Transactions on Multimedia Computing, Communications and Applications*, 20(9): 1–24.
- Xiao, R.; Kim, S.; Georgescu, M.-I.; Akata, Z.; and Alaniz, S. 2025. FLAIR: VLM with Fine-grained Language-informed Image Representations. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*.
- Zawacki, A.; Wu, C.; Shih, G.; Elliott, J.; Fomitchev, M.; Hussain, M.; ParasLakhani; Culliton, P.; and Bao, S. 2019. SIIM-ACR Pneumothorax Segmentation. <https://kaggle.com/competitions/siim-acr-pneumothorax-segmentation>. Kaggle.
- Zhai, X.; Mustafa, B.; Kolesnikov, A.; and Beyer, L. 2023. Sigmoid loss for language image pre-training. In *Proceedings of the IEEE/CVF international conference on computer vision*, 11975–11986.
- Zhang, X.; Wu, C.; Zhang, Y.; Xie, W.; and Wang, Y. 2023. Knowledge-enhanced visual-language pre-training on chest radiology images. *Nature Communications*, 14(1): 4542.
- Zhang, Y.; Jiang, H.; Miura, Y.; Manning, C. D.; and Langlotz, C. P. 2022. Contrastive learning of medical visual representations from paired images and text. In *Machine learning for healthcare conference*, 2–25. PMLR.
- Zhang, Z.; Yu, Y.; Chen, Y.; Yang, X.; and Yeo, S. Y. 2025. Medunifier: Unifying vision-and-language pre-training on medical data with vision generation task using discrete visual representations. *arXiv preprint arXiv:2503.01019*.
- Zhongyi Shui, W. C. S. W. R. G. L. L. L. Y. X. Y. T. L. Q. Z. L. Z., Jianpeng Zhang. 2025. Large-scale and fine-grained vision-language pre-training for enhanced CT image understanding. In *The Thirteenth International Conference on Learning Representations*.