
G2D: From Global to Dense Radiography Representation Learning via Vision-Language Pre-training

Che Liu^{1,2}, Cheng Ouyang^{3,8,9}, Sibò Cheng¹⁰
Anand Shah^{6,7}, Wenjia Bai^{2,3,4}, Rossella Arcucci^{1,2}

¹Department of Earth Science and Engineering, Imperial College London, UK

²Data Science Institute, Imperial College London, UK

³Department of Computing, Imperial College London, UK

⁴Department of Brain Sciences, Imperial College London, UK

⁶Department of Infectious Disease Epidemiology, Imperial College London, UK

⁷Royal Brompton and Harefield Hospitals, UK

⁸Department of Engineering Science, University of Oxford, Oxford, UK

⁹Institute of Clinical Sciences, Imperial College London, UK

¹⁰CEREA, École des Ponts and EDF R&D, Île-de-France, France.

che.liu21@imperial.ac.uk

Abstract

Medical imaging tasks require an understanding of subtle and localized visual features due to the inherently detailed and area-specific nature of pathological patterns, which are crucial for clinical diagnosis. Although recent advances in medical vision-language pre-training (VLP) enable models to learn clinically relevant visual features by leveraging both medical images and their associated radiology reports, current medical VLP methods primarily focus on aligning images with entire reports. This focus hinders the learning of dense (pixel-level) visual features and is suboptimal for dense prediction tasks (e.g., medical image segmentation). To address this challenge, we propose a novel medical VLP framework, named **Global to Dense** level representation learning (**G2D**), which aims to learn global and dense visual features simultaneously using only image-text pairs without extra annotations. In particular, G2D designs a **Pseudo Segmentation (PS)** task, which enables the model to learn dense visual features during VLP. Notably, generating PS masks can be performed on the fly during VLP, which does not incur extra trainable parameters. With this simple yet effective idea, G2D achieves superior performance across 5 medical imaging tasks and 25 diseases. Particularly, in the segmentation task which requires dense visual features, G2D surpasses existing models even with just 1% of the training data for finetuning, compared to 100% used by other models. The code can be found in <https://github.com/cheliu-computation/G2D-NeurIPS24/tree/main>.

1 Introduction

In medical image analysis, learning global and dense visual representations typically requires labor-intensive and costly image and pixel-level annotations [1, 2]. Vision-language pre-training (VLP) attempts addressing this by aligning vision and language content using paired datasets [3, 4, 5, 6]. Although existing medical VLP methods excel at learning global visual features [7], they face challenges with dense visual features because the level of detail in text reports does not offer sufficient

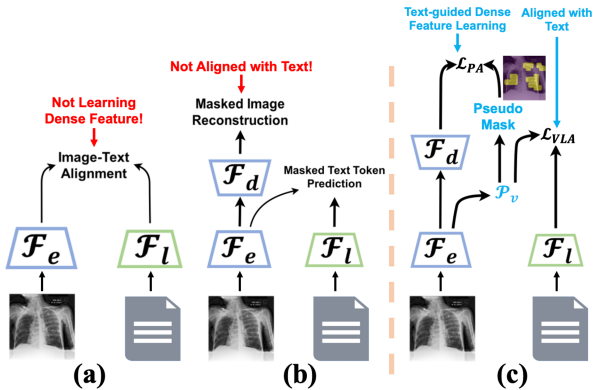


Figure 1: Comparing existing medical VLP methods with G2D: **a)** Alignment-based approaches lack dense (pixel-level) feature learning. **b)** Reconstruction-based approaches do not align with text, resulting in a deficiency in discriminative and clinically relevant visual features. **c)** The framework of **G2D (proposed)** learns dense, clinically relevant, text-aligned visual features through derived pseudo masks and image-text alignment. We use **red** text to highlight the deficiencies of existing methods and **blue** text to emphasize our advantages.

pixel-level supervision for learning these more detailed aspects. Existing medical VLP methods are categorized into two main types, as shown in Fig. 1:

- **Alignment-based Approaches**, which focus on aligning images with reports [4, 8, 9, 5, 6, 2, 10]. Although methods like [4, 8, 9] align images with entire reports and text tokens, they struggle to learn dense, clinically relevant visual features. This is due to the ambiguous supervision targets provided by text tokens, which lack explicit relational pairing with image regions, as discussed in [2].
- **Reconstruction-based Approaches**, which learn representations by reconstructing masked images or reports using masked modeling techniques [11, 12]. However, they also lack success in capturing dense, clinically relevant visual features, as the reconstruction task primarily focuses on low-level patterns (texture, shape) rather than high-level semantics [13].

Despite advancements in medical VLP, limitations still exist. Current alignment approaches align image patches with text tokens in a brute-force manner and possibly cause misalignments when some word tokens (e.g., ‘compatible’ or ‘unremarkable’) lack direct visual counterparts, leading to ambiguous local alignments. Meanwhile, reconstruction-based approaches may ignore high-level image semantics. They are designed to recover low-level visual information such as intensity and texture, without accounting for high-level semantics [14, 13, 15]. As a result, both approaches perform suboptimally for downstream tasks, such as semantic segmentation and visual grounding, which require learning of granular visual features that are aligned with high-level semantics.

While numerous VLP methods are designed to capture dense visual features for natural image datasets (e.g., ImageNet) they often struggle to transfer directly to medical images because they depend on a well-trained object detection model [16, 17] or a well-aligned VLP model [18, 19]. In the medical domain, obtaining such pre-trained models is difficult as objects can be defined in various ways within a single medical image (e.g., based on organs, anatomical structures, or abnormal regions). Additionally, in medical domain, there is a lack of foundational VLP models that are both publicly accessible and are trained on sufficiently large image-text pairs that cover diverse medical imaging applications.

In response to the aforementioned challenges, we introduce a novel medical VLP approach termed G2D. This approach is designed to extract global and dense visual representations from radiography along with their associated radiology reports, with improved **feature granularity** and enriched **semantic information**. Central to our approach is a pretext task, **Pseudo Segmentation (PS)**, which is guided by a pseudo mask (segmentation target) derived from a carefully refined and filtered attention map. PS encourages the model to learn dense representations through a pixel-level pretext task that incorporates high-level semantics. This approach, in contrast to traditional methods that align image patches with text tokens, inherently mitigates the misalignment bias and allows learning of more representative features. Notably, the PS pretext task can be implemented to run concurrently with vision-language alignment, ensuring that the model can be trained end-to-end, contrasting with the two-stage training methods [18].

To evaluate the effectiveness of G2D relative to other state-of-the-art (SOTA) VLP approaches, we deploy the pre-trained model across a diverse range of downstream tasks, including medical image classification, semantic segmentation, object detection, as well as zero-shot image classification and

visual grounding, on six public large-scale CXR datasets. The experimental results demonstrate the superior performance of G2D over existing VLP approaches on these medical applications. Overall, our contribution is three-fold:

1. We introduce G2D, the first end-to-end encoder-decoder medical VLP approach designed to learn visual representations from the global level down to the dense level, supervised by paired radiology reports and a pixel-wise pretext task.
2. We carefully design a pretext task tailored for medical VLP, pseudo segmentation. It formulates a pseudo mask as segmentation target, allowing the model to learn dense visual representations in the pretext task which can benefit downstream dense visual tasks in medicine. The pseudo mask can be generated using a parameter-free processor that leverages the attention map derived from the visual representation associated with radiology reports.
3. We conduct comprehensive experiments to validate the efficacy of the proposed G2D approach, which outperforms peer approaches across five uni-modal and cross-modal downstream tasks.

2 Related Works

Alignment-based Medical VLP. Drawing inspiration from [3], aligning images with their corresponding textual descriptions in the latent space has led to notable advancements in VLP. Within the CXR domain, while ConVIRT [4] made an early attempt at employing bidirectional contrastive learning to globally align entire images with their paired reports, there remained room for refinement. GLoRIA [8] and MGCA [9] represent advancements in image-report alignment, introducing sophisticated global-local methodologies to the field [8, 9]. These approaches endeavor to establish correspondences between distinct image and text tokens. However, it is crucial to recognize that the granularity of token-level alignment could inadvertently introduce distortions to the medical context, potentially leading to misalignments, as illustrated by [20, 2]. Med-UniC [20] utilizes augmented text in VLP training to cultivate language invariance, with the goal of mitigating linguistic biases from VLP. Meanwhile, MedKLIP [5] and KAD [21] harness domain-specific knowledge from external annotated datasets to enhance textual information extraction. Notably, these approaches [20, 5, 21] are contingent upon external resources or extra data to optimize cross-modal representation learning, which could potentially constrain their generalizability.

Reconstruction-based Medical VLP. Several studies, including [12, 11, 22], have employed reconstruction of image and text tokens as a pretext task within VLP. Specifically, MRM [12] endeavors to reconstruct the original image from a masked version and simultaneously aims to regenerate the original text using both the masked image and text as inputs. Conversely, PRIOR [11] adopts a strategy that focuses on cross-modal representation by reconstructing images and sentences based on complete image and report inputs. An enhancement to the MRM [12] approach is proposed by [22], where token weights are adjusted during the reconstruction phase.

While these methods have demonstrated promising outcomes, the ability of the reconstruction pretext task to capture high-level semantic representations is limited, as shown in [14, 15, 13], and is further challenged by the absence of explicit semantic-related constraints in dense visual representation learning.

3 Methodology

The central aim of G2D is to learn global and dense visual representations from medical images under the supervision of their corresponding radiology reports. As illustrated in Fig 2 Left, G2D integrates two alignment strategies: vision-language alignment (VLA) that learns global representations, and pixel alignment (PA) that focuses on granular representation via a pixel-level pretext task, **Pseudo Segmentation (PS)**. The pseudo mask for PS is constructed through a parameter-free mechanism, which is operated alongside VLA. The PS pretext task enables G2D to derive dense representations at both encoder and decoder levels during pre-training. Moreover, the task head of the pretext task facilitates a smoother transfer for the pre-trained encoder to be applied to downstream segmentation tasks, reducing the gap between the dense visual representation learned from VLP and the needs of downstream dense visual tasks after VLP. This contrasts with previous meth-

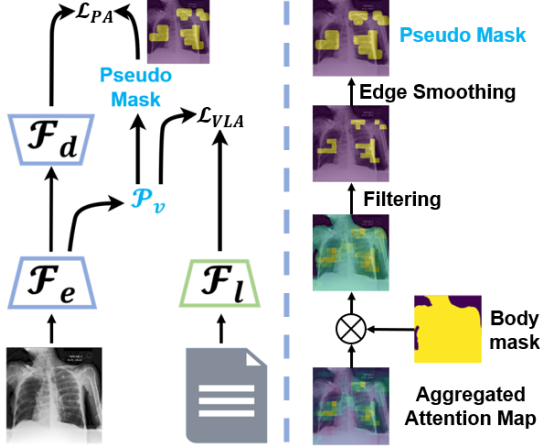


Figure 2: **Left:** Framework of G2D. **Right:** Pipeline for pseudo mask construction. We visualize the constructed pseudo mask and corresponding sentence in the radiology report in Sec A.7.

ods [4, 8, 9, 21, 5, 6] that typically transfer only the pre-trained encoder, potentially leading to an information gap between the pre-training and downstream tasks.

3.1 Vision-Language Contrastive Learning

We utilise a dual-encoder image-text contrastive approach following [4, 8, 9, 5]. Given a training set S consisting of N pairs of image-text (v_i, l_i) , where $v_i \in \mathcal{V}$ denotes an image and $l_i \in \mathcal{L}$ denotes a text report, $i = 1, 2, 3, \dots, N$, G2D employs an image encoder $\mathcal{F}_e : \mathcal{V} \mapsto \mathbb{R}^{D_v}$ to encode the image into an embedding of dimension D_v , and a text encoder $\mathcal{F}_l : \mathcal{L} \mapsto \mathbb{R}^{D_l}$ to encode the text report into an embedding of dimension D_l . The embedded image and text features can be denoted as $\mathbf{S} = \{(\mathbf{v}_1, \mathbf{l}_1), (\mathbf{v}_2, \mathbf{l}_2), \dots, (\mathbf{v}_N, \mathbf{l}_N)\}$, where $\mathbf{v}_i = \mathcal{F}_e(v_i)$ and $\mathbf{l}_i = \mathcal{F}_l(l_i)$.

As depicted in Fig. 2, G2D incorporates two alignment strategies: VLA and PA. For VLA, the model aims to learn global visual and text representations by pulling the embeddings of paired image-report samples closer, while distancing embeddings of unpaired samples, using a contrastive loss \mathcal{L}_{VLA} . The objective of contrastive learning is to predict N positive matched pairs (v_i, l_i) and $N^2 - N$ negative pairs among $N \times N$ possible image-text pair combinations [3]. Subsequently, two non-linear vision and language projectors \mathcal{P}_v and \mathcal{P}_l transform \mathbf{v}_i and \mathbf{l}_i into the same dimension d , where $\hat{\mathbf{v}}_i = \mathcal{P}_v(\mathbf{v}_i)$, $\hat{\mathbf{l}}_i = \mathcal{P}_l(\mathbf{l}_i)$, and $\hat{\mathbf{v}}_i, \hat{\mathbf{l}}_i \in \mathbb{R}^d$. After obtaining image feature vectors $[\hat{\mathbf{v}}_i]_{i=1}^N$ and text feature vectors $[\hat{\mathbf{l}}_i]_{i=1}^N$ with the same dimension d , the contrastive loss \mathcal{L}_{VLA} can be formulated as:

$$\mathcal{L}_{VLA} = -\frac{1}{K} \sum_{i=1}^N \left(\log \frac{\exp(\hat{\mathbf{v}}_i^\top \hat{\mathbf{l}}_i / \sigma)}{\sum_{j=1}^K \exp(\hat{\mathbf{v}}_i^\top \hat{\mathbf{l}}_j / \sigma)} \right) \quad (1)$$

σ denotes the temperature hyper-parameter empirically set to 0.07 following [9], and $K \in N$ is the batch size.

3.2 Pseudo Segmentation Mask Construction

Notably, although MedSAM [23] claims to build image-mask pairs, it requires box prompt inputs not available in the MIMIC-CXR [24] dataset. Designing a box prompt for each image is labor-intensive and unfeasible for this work, so we construct the pseudo mask based on attention maps.

Attention Aggregation. Inspired by CLIP [3], we incorporate an attention pooling mechanism in conjunction with the non-linear projector \mathcal{P}_v to derive a pixel-wise attention map. A dense feature map \mathbf{V}_i is extracted from the final convolutional layer before the pooling operation in the image encoder \mathcal{F}_e , with the dimension $C \times H \times W$. Here, C denotes the number of channels, while H and W represent the height and width of the feature maps. Subsequently, we reshape \mathbf{V}_i into a dimension of $HW \times C$. In this way, \mathbf{V}_i can be interpreted as a sequence of pixel embeddings, where each token in this sequence represents the embedding of an individual pixel. The length of this sequence is defined by the number of channels, C . A special token, [CLS], is introduced to aggregate all pixel embeddings through multi-head self-attention (MHSA) [25, 3]. This process offers an attention score matrix W_i^h for each pixel, with dimensions $h \times H \times W$. Here, h signifies the

attention head number, and $h \in \mathcal{H}$, with \mathcal{H} being the total number of attention heads. This attention score matrix characterizes the information exchange between pixels and semantics provided by the text [3, 18], and therefore it carries semantic information and is an ideal candidate for constructing the pretext pseudo mask. To derive the pseudo mask, we aggregate W_i^h across all attention heads to produce \hat{W}_i , as described by:

$$\hat{W}_i = \frac{\sum_{h=1}^{\mathcal{H}} W_i^h}{h} \quad (2)$$

Mask Filtering and Edge Smoothing. After obtaining the aggregated attention map \hat{W}_i , we up-sample it to match the original image dimensions $H' \times W'$. To remove pseudo mask regions in the background, we construct a body mask for each CXR image using a histogram-based thresholding approach, following common practice [26, 27]. Subsequently, all attention scores outside the body mask are set to zero. A threshold is applied to filter out low attention scores within the body mask, transforming \hat{W}_i into a binary mask. The threshold is determined at the 85% percentile of attention scores from \hat{W}_i . The binary pseudo mask M_i is formulated as:

$$M_i^{j,k} = \begin{cases} 1 & \text{if } W_i^{j,k} \geq \text{threshold} \\ 0 & \text{otherwise} \end{cases}, \text{ where} \\ j = 1, 2, 3, \dots, H', k = 1, 2, 3, \dots, W' \quad (3)$$

To smooth the square-like boundary in the mask caused by upsampling, we apply bilateral filtering (BF) [28] to M_i , resulting in a refined pseudo mask \tilde{M}_i , as shown in Fig. 2 Right. A comprehensive ablation study discussing the threshold and smoothing operation is presented in Sec. 4.5.

3.3 Dense Visual Representation Learning through Pseudo Segmentation in VLP

While the global visual representation can be learned via VLA, dense representation often lacks direct alignment. To tackle this limitation, we introduce an image decoder, denoted as \mathcal{F}_d , as shown in Fig. 2 Left. This decoder takes visual feature \mathbf{V}_i as input and utilises the pseudo mask \tilde{M}_i as the supervisory signal for the pretext task. We employ the commonly used soft Dice loss and binary cross-entropy loss [27] to optimise this task. The training loss function for \mathcal{L}_{PA} is formulated as:

$$\begin{aligned} \mathcal{L}_{\text{PA}} &= \frac{1}{2}(\mathcal{L}_{\text{Dice}} + \mathcal{L}_{\text{BCE}}), \\ \mathcal{L}_{\text{Dice}} &= \sum_{i=1}^K \sum_{j=1}^{H'} \sum_{k=1}^{W'} \left(1 - \frac{2 \times (\tilde{M}'_{i,j,k} \odot \tilde{M}_{i,j,k})}{\tilde{M}'_{i,j,k} + \tilde{M}_{i,j,k}} \right), \\ \mathcal{L}_{\text{BCE}} &= - \sum_{i=1}^K \sum_{j=1}^{H'} \sum_{k=1}^{W'} \left[\tilde{M}_{i,j,k} \log(\tilde{M}'_{i,j,k}) + (1 - \tilde{M}_{i,j,k}) \log(1 - \tilde{M}'_{i,j,k}) \right], \\ &\text{with } \tilde{M}'_i = \mathcal{F}_d(\mathbf{V}_i) \end{aligned} \quad (4)$$

The total loss for G2D is the sum of the VLA loss (Eq. 1) and the PA loss (Eq. 4):

$$\mathcal{L}_{\text{total}} = \mathcal{L}_{\text{VLA}} + \mathcal{L}_{\text{PA}} \quad (5)$$

It is worth noting that the pseudo mask is designed as a pixel-wise pretext supervisory signal. Although there is no manual annotation involved, the pseudo mask is constructed from the visual feature of the image encoder, which is pre-trained to align with radiology reports and thus contains clinical knowledge such as anatomical regions mentioned by the reports. In this sense, it can be a good surrogate target for learning pixel-wise semantic information. To demonstrate that the pseudo mask serves as a meaningful target for dense visual pre-training, we conduct an ablation study to use a perturbed pseudo mask with corrupt semantics for pre-training, and compare it to the proposed pseudo mask, as detailed in Table 8 and Sec A.6.

4 Experiments and Analysis

In this section, we compare our approach with SOTA medical VLP techniques. The implementation details and dataset training/test splits are reported in Sec A.3, A.4.

Pretraining Dataset and Configuration We utilise the MIMIC-CXR dataset [29, 24]. After preprocessing based on established protocols [9, 5], it provides 213,384 image-text pairs for pre-training. For the VLP part, we employ a standard ResNet-50 as the vision encoder \mathcal{F}_e and adopt the decoder part of a U-Net as the vision decoder \mathcal{F}_d . We adopt ClinicalBERT [30] as the text encoder using configurations described in [5, 21]. In line with [9, 8], G2D is pre-trained for 50 epochs across 16 A100 GPUs, each accommodating a batch size of 128. The AdamW optimizer is employed with a learning rate set to 2×10^{-4} and a weight decay of 1×10^{-8} . Additionally, a linear warm-up and a cosine annealing scheduler are incorporated in the training process.

4.1 Downstream Task Datasets and Configurations

For downstream tasks, our focus is to evaluate the efficacy of G2D in learning granular visual features that can be used for localisation, vision-language understanding, and visual recognition tasks. We examine the capability and transferability of the learned cross-modal representations by using them for five distinct medical imaging tasks, covering a spectrum of 25 different diseases.

Medical Image Segmentation. This task utilises the RSNA [31] and SIIM [32] datasets, following preprocessing guidelines established in [9, 8]. We adopt U-Net [1] fine-tuning configurations following [8, 9]. The pre-trained vision encoder is frozen, while only the decoder parameters are updated during fine-tuning. Performance is assessed using the Dice score, following the evaluation protocol in [8, 9]. It is noteworthy that the original MedKLIP [5] uses a different configuration (*updating the vision encoder*) compared to other methods (*freezing the vision encoder*) [4, 8, 9, 6]. Therefore, in these experiments, we reference the results reported in [20], which reimplemented MedKLIP under a setting consistent with all other methods. For a fair comparison specifically with MedKLIP, we also reimplement G2D under MedKLIP’s original setting, as reported in the Sec A.5.

Medical Object Detection. This task is conducted using the RSNA dataset [31] for Pneumonia Detection and the Object-CXR dataset [33] for Foreign Objects Detection, adhering to preprocessing methods from [9]. We employ YOLOv3 [34] for detection, using the pre-trained vision encoder and updating an additional detection head during fine-tuning. We report the mean Average Precision (mAP) with IoU thresholds between 0.4~0.75. The setup for this task is in accordance with in [9].

Zero-shot Medical Image Visual Grounding. In accordance with [5], this task is conducted on the RSNA [31] and SIIM [32] datasets, using the same official data split and evaluation metrics. We employ CXR images as input and utilise the corresponding ground truth label maps for assessing the grounding performance, in terms of recall, IoU, and Dice score.

Zero-shot Medical Image Classification. In compliance with the guidelines set forth in [5, 21], we conduct this task on the RSNA [31], SIIM [32], CheXpert [35], and CXR14 [36] datasets. For the RSNA and SIIM datasets, we employ the test set splits provided by MedKLIP [5], given that KAD [21] did not conduct experiments on these two datasets. For the CheXpert and CXR14 datasets [35, 36], we use the official test set splits to ensure a fair comparison with KAD [21]. It is important to note that MedKLIP [5] creates its own test split rather than using the official test split. Hence, we do not use MedKLIP’s splits in our experiments. We report the results using the macro average of AUC, F1, and ACC scores across all diseases.

Medical Image Fine-tuned Classification. In alignment with [5, 21], we use the CXR14 dataset [36], comprising 112,120 frontal-view X-rays from 30,805 patients, annotated for 14 diseases. We adhere to the official split for consistent evaluation, following KAD [21]. It is worth noting that MedKLIP does not use the official data split. Hence, we refer to the results reported in KAD [21] rather than those from the original MedKLIP [5]. To ensure a fair comparison with MedKLIP, we reimplemented G2D for this experiment under the MedKLIP configuration, as detailed in Sec A.5. CXR images are resized to 256×256 [21]. During fine-tuning, all model parameters are updated, including the pre-trained vision encoder and linear classifier. The AdamW optimizer is used with a learning rate of 1×10^{-4} and a batch size of 64 for 50 epochs. Evaluation is based on the AUC score, adhering to the protocol outlined in [8, 9, 12].

Table 1: Results of semantic segmentation and object detection. Best results are highlighted in bold, with ‘-’ denoting mAP values < 1%. Methods with * use disease-level annotations. ‘/’ indicates object detection not deployable with encoder-decoder architecture. The MedKLIP results in this table differ from the original work [5] because MedKLIP fine-tuned the encoder in their original study, whereas other methods froze the encoder. To ensure fairness, we reimplemented MedKLIP with the frozen encoder for comparison in this table. Additionally, for a fair comparison specifically with MedKLIP, we compare G2D with MedKLIP under its original configuration in Tab 7 and Sec A.5.

Tasks	Semantic Segmentation (Dice)						Object Detection (mAP)					
	SIIM			RSNA			RSNA			Object CXR		
Datasets	1%	10%	100%	1%	10%	100%	1%	10%	100%	1%	10%	100%
Random Init	9.0	28.6	54.3	6.9	10.6	18.5	1.0	4.0	8.9	-	0.5	4.4
ImageNet Init	10.2	35.5	63.5	34.8	39.9	64.0	3.6	8.0	15.7	-	2.9	8.3
ConVIRT [4]	25.0	43.2	59.9	55.0	67.4	67.5	8.2	15.6	17.9	-	8.6	15.9
GLoRA [8]	35.8	46.9	63.4	59.3	67.5	67.8	9.8	14.8	18.8	-	10.6	15.6
GLoRIA-MIMIC [8]	37.4	57.1	64.0	60.3	68.7	68.3	11.6	16.1	24.8	-	8.90	16.6
MGCA [9]	49.7	59.3	64.2	63.0	68.3	69.8	12.9	16.8	24.9	-	12.1	19.2
M-FLAG [6]	52.5	61.2	64.8	64.6	69.7	70.5	13.7	17.5	25.4	-	12.4	19.3
MedKLIP* [5]	50.2	60.8	63.9	66.2	69.4	71.9	8.9	16.3	24.5	-	7.1	11.6
Ours (encoder)	62.6	63.1	66.8	70.9	72.6	75.1	15.9	21.7	27.2	3.8	13.1	20.4
Ours (encoder-decoder)	65.6	66.9	68.4	72.8	73.4	76.9	/					

Medical Image Linear Classification. In strict accordance with the configuration in [8, 4, 9], this task is conducted on the CheXpert [35], RSNA [31], and COVIDx [37] datasets. We only update a randomly initialized linear classification layer, while the pre-trained vision encoder remains frozen. For fair evaluation, we employ AUC scores on CheXpert and RSNA, along with accuracy metrics on COVIDx, as mentioned in [8, 9]. Apart from zero-shot image classification and visual grounding, we fine-tune using 1%, 10%, 100% of the training data for all downstream tasks. Detailed settings, including implementation and data splits, are outlined in Sec A.4.

4.2 Performance on Visual Localisation Tasks

In Tab 1, following [16, 38], we evaluate G2D alongside other SOTA approaches on two pivotal visual localisation tasks: semantic segmentation and object detection. The aim is to assess the efficacy of the dense visual features learned.

Initially, we transfer only the encoder weights from the pre-trained G2D for the segmentation task, adhering to the protocols of [9, 8, 4, 6]. In this setup, our approach consistently achieves the highest performance across all data fractions for both SIIM [32] and RSNA datasets [31]. To assess the impact of the visual decoder pre-trained with the PS pretext task, we transfer the weights of both the encoder and decoder from G2D for the segmentation task, resulting in striking outcomes. Remarkably, with just 1% of training data, G2D surpasses the performance of all peer methods, even those trained with a full 100% of training data. This observation underlines the fact that the pixel-level pretext task, PS, significantly improves the quality of dense visual features derived from VLP, which provide advantages for the downstream segmentation task.

In object detection, our method consistently outperforms existing methods across all data fractions for both RSNA and Object-CXR datasets [31, 33]. Notably, G2D achieves a 3.8% mAP on the Object-CXR dataset with just 1% of the data for fine-tuning, a significant leap from other methods that scarcely reach a 1% mAP.

These results highlight the efficacy of our proposed model, G2D, and the pretext task, PS, especially in semantic segmentation tasks that rely on dense visual features. PS not only enables G2D to learn visual representations in the encoder-decoder structure but also reduces the gap between pre-training and downstream tasks. By enhancing the encoder’s ability to capture global and dense features simultaneously, PS surpasses existing approaches, proving particularly advantageous for object detection tasks that heavily rely on dense features [39].

4.3 Performance on Vision-Language Understanding

In Tab 2, we evaluate the efficacy of G2D on vision-language understanding tasks, zero-shot visual grounding and zero-shot image classification. For the zero-shot visual grounding task, our proposed method outperforms peer approaches. Specifically, on the SIIM dataset [32], it achieves a leading Dice score of 5.1. This dominance persists in the RSNA dataset [31], where our method reaches a

Table 2: Comparison between G2D (ours) and various other medical VLP methods in vision-language understanding tasks, with the best results emphasized in bold. Methods marked with \star utilize extra annotated data during pre-training. ‘/’ indicates that the original work did not report the results. Notably, KAD [21] does not report ACC for the CheXpert dataset.

(a) Results of zero-shot visual grounding task.

Task		Zero-shot Visual Grounding					
Datasets Methods	SIIM			RSNA			
	Recall	IoU	Dice	Recall	IoU	Dice	
GLoRIA [8]	23.8	1.2	2.1	83.3	21.8	34.7	
BioViL [40]	19.6	1.7	2.6	85.2	30.3	43.9	
MedKLIP \star [5]	35.6	2.1	4.0	86.6	31.7	46.5	
Ours	37.7	3.9	5.1	88.4	33.5	47.7	

(b) Results of zero-shot image classification task.

Task		Zero-shot Image Classification										
Datasets Methods	AUC	RSNA			SIIM			CXR14			CheXpert	
		F1	ACC	AUC	F1	ACC	AUC	F1	ACC	AUC	F1	
ConVIRT [4]	80.4	58.4	76.1	64.3	43.3	57.0	56.0	13.5	45.9	59.0	26.4	
GLoRIA [8]	71.5	49.0	71.3	53.4	38.2	40.5	61.0	17.4	50.3	75.0	57.0	
BioViL [40]	82.8	58.3	76.7	70.8	48.6	69.1	66.2	66.2	63.3	69.3	46.3	
CheXzero \star [5]	85.8	62.1	79.4	68.8	47.0	54.7	/	/	/	88.9	60.6	
MedKLIP \star [5]	86.9	63.4	80.0	89.2	68.3	84.3	72.6	24.4	79.6	87.9	61.4	
KAD \star	/	/	/	/	/	/	78.9	32.3	81.6	90.5	64.6	
Ours	87.6	64.8	81.5	89.7	69.3	85.4	79.4	33.1	82.3	91.2	65.6	

Table 3: Evaluation of image classification fine-tuning on the CXR14 dataset is conducted, with all metrics presented as AUC scores, where the mean metric is macro-averaged. Best performances are highlighted in bold. Methods marked with \star utilize extra annotated data for pre-training. MedKLIP’s results here differ from the original study [5] as it did not utilize the official test split, unlike KAD [21]. We use the result of MedKLIP reported by KAD [21], which reimplemented MedKLIP on the official test set for fairness. All results in this table are sourced from KAD [21]. To compare fairly with MedKLIP, we assess G2D against its original configuration in Tab 7 and Sec A.5.

Data Fraction	Method	Mean	Atelectasis	Cardiomegaly	Effusion	Infiltration	Mass	Nodule	Pneumonia	Pneumothorax	Consolidation	Edema	Emphysema	Fibrosis	Pleural Thickening	Hernia
1%	Random Init	58.1	55.7	57.7	63.6	61.6	55.0	60.2	57.1	58.2	60.8	63.3	53.4	63.7	56.8	46.0
	ImageNet Init	63.5	66.2	64.2	72.1	57.0	59.0	58.5	60.0	62.6	62.4	66.8	61.5	70.7	63.1	64.5
	ConVIRT [4]	64.9	66.0	78.2	78.9	61.1	59.6	65.5	60.8	68.8	65.7	60.7	65.8	68.0	62.7	46.6
	GLoRIA [8]	59.7	59.7	56.7	74.1	64.6	55.9	55.7	61.1	60.7	66.5	66.9	55.0	55.8	59.2	43.6
	BioViL [40]	57.9	55.5	56.4	72.2	65.0	56.7	54.6	62.6	56.0	65.7	68.1	51.6	51.3	59.2	36.0
	MedKLIP \star [5]	60.9	65.5	59.0	74.5	64.3	55.0	61.1	60.9	59.9	65.9	68.2	53.5	64.8	59.3	40.0
	KAD \star [21]	78.7	77.0	88.2	82.9	69.2	75.1	69.7	73.5	86.1	72.7	81.3	89.3	74.3	69.2	93.8
	Ours	79.1	78.1	88.3	83.1	70.2	75.4	69.7	74.0	86.5	72.9	81.6	90.2	74.4	69.5	94.1
10%	Random Init	69.1	68.2	76.6	74.6	67.4	62.3	58.0	63.6	72.8	67.8	78.0	64.7	71.5	65.3	77.1
	ImageNet Init	72.6	70.9	79.8	76.9	68.4	69.3	65.6	63.0	79.3	67.1	76.7	74.9	72.9	71.1	81.0
	ConVIRT [4]	77.1	74.0	84.3	81.1	69.3	74.8	70.0	67.1	82.8	70.1	81.4	87.1	76.7	71.9	89.3
	GLoRIA [8]	74.3	72.1	80.8	80.0	68.7	73.3	67.5	65.8	77.9	67.6	79.7	79.9	78.7	69.3	78.7
	BioViL [40]	72.7	70.3	78.5	79.0	66.6	71.8	67.1	66.5	76.7	68.4	79.9	76.1	74.8	65.3	76.3
	MedKLIP \star [5]	74.8	72.9	80.2	79.3	69.8	71.9	68.1	66.6	79.6	69.6	81.1	79.5	75.6	71.3	81.9
	KAD \star [21]	80.7	77.6	88.9	83.3	71.8	78.3	71.9	73.7	87.2	75.0	83.3	90.3	80.7	72.3	95.3
	Ours	81.1	78.4	89.3	83.7	72.2	78.8	72.3	74.1	87.8	75.3	84.0	90.4	80.8	72.5	95.4
100%	Random Init	79.0	75.0	87.9	81.5	69.1	79.8	72.6	70.3	82.6	73.1	83.9	83.5	80.7	75.4	90.3
	ImageNet Init	80.4	76.3	86.7	82.3	69.3	82.3	76.3	71.9	84.0	73.7	84.2	89.3	81.9	77.0	89.9
	ConVIRT [4]	80.8	77.1	86.7	82.5	70.3	81.8	76.1	72.2	85.7	74.7	85.4	90.1	80.9	77.1	90.9
	GLoRIA [8]	80.0	76.0	85.5	81.8	70.0	81.4	74.9	71.5	82.8	73.9	83.2	88.7	81.3	76.7	92.1
	BioViL [40]	80.0	76.5	87.1	82.4	69.7	81.9	75.2	71.0	84.5	74.2	84.2	87.1	82.1	75.9	88.8
	MedKLIP \star [5]	80.1	76.4	84.9	82.3	69.7	82.0	74.7	71.2	83.9	75.1	84.8	87.9	81.7	77.7	89.2
	KAD \star [21]	82.5	78.5	89.7	84.0	71.3	83.6	77.1	74.0	87.4	75.3	86.0	91.6	82.9	77.8	96.1
	Ours	83.1	79.9	90.2	84.5	71.8	84.2	78.0	74.2	87.7	75.6	86.9	92.0	83.1	78.2	96.5

Dice score of 47.7, surpassing other SOTA approaches. When examining zero-shot image classification, our method again shows its superiority across the AUC, F1, and ACC metrics on both the RSNA [31] and SIIM datasets [32]. Such consistent and superior outcomes underscore the adaptability and effectiveness of G2D in handling vision-language understanding tasks, indicating that integrating PS into G2D can enhance not only uni-modal but also cross-modal tasks.

4.4 Performance on Visual Recognition Tasks

In our final assessment focused on visual recognition, Tab 3 demonstrates our method’s consistent supremacy on the CXR14 dataset [36] for fine-tuned disease classification across 1%, 10%, and 100% training data. Similarly, Tab 4 underscores that G2D achieves the highest performance on the CheXpert, RSNA, and COVIDx datasets [35, 31, 37] for linear evaluation across all training data ratio. Notably, G2D consistently outperforms even those methods like MedKLIP and KAD [41] that leverage additional disease-level annotations during pre-training stage. This demonstrates G2D’s representative visual features, suggesting that enhancing dense representation learning via PS can also improve results in tasks primarily anchored on global representation.

Table 4: Linear classification results for CheXpert, RSNA, and COVIDx datasets with 1%, 10%, and 100% training data. The best results are highlighted in bold. Methods with \star leverage disease-level annotations for pre-training. The evaluation metric follows [9].

Datasets (Metric) Methods	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 [4]	85.9	86.8	87.3	77.4	80.1	81.3	72.5	82.5	92.0
GLoRIA [8]	86.6	87.8	88.1	86.1	88.0	88.6	67.3	77.8	89.0
GLoRIA-MIMIC [8]	87.1	88.7	88.0	87.0	89.4	90.2	66.5	80.5	88.8
MGCA [9]	87.6	88.0	88.2	88.6	89.1	89.9	72.0	83.5	90.5
MRM [12]	88.5	88.5	88.7	91.3	92.7	93.3	66.9	79.3	90.8
MedKLIP \star [5]	86.2	86.5	87.7	87.3	88.0	89.3	74.5	85.2	90.3
Ours	89.7	90.4	91.1	92.2	92.9	93.6	76.6	88.2	93.4

Table 5: Results of various ablation experiments. The best results are bolded.

(a) Loss for the decoder. ‘None’ indicates Encoder-Only visual backbones. (b) Threshold for constructing pseudo segmentation masks. (c) Ablation of the number of dimensions of projectors.

Decoder Loss	SIIM Dice	RSNA mAP	CXR14 AUC	Threshold	SIIM Dice	RSNA mAP	CXR14 AUC	Num of Dim	SIIM Dice	RSNA mAP	CXR14 AUC
None	49.2 \pm 1.5	11.7 \pm 1.2	77.1 \pm 1.5	85% percentile	65.6\pm1.7	15.9\pm0.8	79.1\pm1.2	128	65.6\pm1.7	15.9 \pm 0.8	79.1\pm1.2
Reconstruction	53.4 \pm 1.3	13.0 \pm 0.9	77.3 \pm 2.1	75% percentile	63.0 \pm 2.1	14.1 \pm 1.2	78.3 \pm 2.0	256	64.9 \pm 1.9	16.1\pm1.1	78.3 \pm 1.5
Pseudo Seg (Ours)	65.6\pm1.7	15.9\pm0.8	79.1\pm1.2	median	58.8 \pm 1.6	12.5 \pm 2.3	75.6 \pm 1.1	512	64.6 \pm 1.2	15.7 \pm 1.0	78.0 \pm 1.3
				GMM [42]	59.2 \pm 1.5	12.9 \pm 1.4	75.2 \pm 1.9				

(d) Ablation of multi-head attention maps aggregation.

Method	SIIM Dice	RSNA mAP	CXR14 AUC
w Aggregation	65.6\pm1.7	15.9\pm0.8	79.1\pm1.2
w/o Aggregation	62.1 \pm 2.2	13.5 \pm 1.7	77.5 \pm 2.3

(e) Number of attention heads.

Heads	SIIM Dice	RSNA mAP	CXR14 AUC
1	63.4 \pm 2.0	14.2 \pm 1.4	78.2 \pm 1.0
2	64.7 \pm 1.6	15.1 \pm 2.3	78.8 \pm 1.5
3	65.6\pm1.7	15.9\pm0.8	79.1\pm1.2
4	65.3 \pm 1.6	15.4 \pm 0.9	78.7 \pm 1.9

(f) Refinement of Pseudo Segmentation Masks

	SIIM Dice	RSNA mAP	CXR14 AUC
w/o body mask	63.4 \pm 1.5	15.3 \pm 2.1	78.4 \pm 1.6
w/o edge smoothing	64.1 \pm 1.2	15.2 \pm 1.7	78.5 \pm 2.2
w both (Ours)	65.6\pm1.7	15.9\pm0.8	79.1\pm1.2

4.5 Ablation Studies

Pseudo Segmentation vs. Reconstruction. In Tab 5a, we evaluate the impact of the proposed PS pretext task in comparison to pixel reconstruction and models without a decoder-level constraint. The model pre-trained with PS outperforms the other two approaches across all three downstream tasks, particularly in semantic segmentation. While the model pre-trained with a pixel reconstruction constraint exhibit improved performance compared to unconstrained variants, such models still underperform the model with the PS constraint. These results underscore the effectiveness of decoder-level pretext tasks and suggest that an emphasis on high-level semantics, derived from PS, is more beneficial than focusing on the low-level semantics from pixel reconstruction. The PS potentially reduces the disparity between features learned through VLP and those required by downstream semantic segmentation tasks. It also enables the model to acquire more representative features that are beneficial for various tasks.

Threshold of Pseudo Mask Construction. As shown in Tab 5b, performance varies with different thresholds, with the 85% percentile threshold proving most effective across all three downstream tasks. Despite employing the Gaussian Mixture Model (GMM) for pseudo mask creation, as suggested by [42], its performance is still surpassed by the 85% percentile approach. This indicates that the original attention map might contain noise, and a higher threshold is beneficial for generating more effective pseudo masks.

Furthermore, Tab 5d highlights the importance of aggregating multi-head attention maps for mask construction. Given the absence of explicit semantic supervision in the PS pretext task, not aggregating these maps leads to the creation of multiple pseudo masks. This excess of masks introduce ambiguous training objectives for VLP.

Impact of Mask Refinement. Refinement of the pseudo masks affects the model’s efficacy, as shown in Tab 5f. Performance tends to decrease when either the body mask is omitted or edge smoothing is not applied. However, integrating both these strategies, as we implement in G2D, yields optimal results. This underscores the vital role of pseudo mask refinement in enhancing model performance.

Ablation on Hyperparameters. We further ablate the number of attention heads and projector dimensionality. Performance improves with more attention heads, peaking at 3 before slightly declin-

ing at 4 (Tab 5e). Optimal segmentation and classification results are achieved with 128-dimensional projectors. While 256 dimensions provide slight benefits for object detection, they reduce performance in other tasks (Tab 5c). Projectors of 512 dimensions do not yield further gains. Thus, we select 3 attention heads and 128-dimensional projectors for an optimal balance of complexity and effectiveness.

5 Conclusion

In this study, we introduce G2D, a novel medical VLP framework for learning global and dense-level representations. Our proposed pixel-level pretext task, pseudo segmentation, leverages a refined attention map to predict a pseudo mask, capturing dense visual features during VLP without requiring additional trainable parameters for its construction. Our model pretrained with this pretext task achieves superior performance across five diverse medical imaging tasks and outperforms methods pretrained with annotated data [5, 21], especially in semantic segmentation. Specifically, on the SIIM [32] dataset, G2D, when fine-tuned with only 1% of the training data, outperforms other medical VLP approaches that utilize the full 100% training set. We anticipate that G2D will inspire further exploration of novel and clinically-guided pretext tasks for medical VLP.

References

- [1] O. Ronneberger, P. Fischer, and T. Brox, “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*. Springer, 2015, pp. 234–241.
- [2] C. Liu, S. Cheng, M. Shi, A. Shah, W. Bai, and R. Arcucci, “Imitate: Clinical prior guided hierarchical vision-language pre-training,” *arXiv preprint arXiv:2310.07355*, 2023.
- [3] A. Radford, J. W. Kim, C. Hallacy, A. Ramesh, G. Goh, S. Agarwal, G. Sastry, A. Askell, P. Mishkin, J. Clark *et al.*, “Learning transferable visual models from natural language supervision,” in *International Conference on Machine Learning*. PMLR, 2021, pp. 8748–8763.
- [4] Y. Zhang, H. Jiang, Y. Miura, C. D. Manning, and C. P. Langlotz, “Contrastive learning of medical visual representations from paired images and text,” *arXiv preprint arXiv:2010.00747*, 2020.
- [5] C. Wu, X. Zhang, Y. Zhang, Y. Wang, and W. Xie, “Medklip: Medical knowledge enhanced language-image pre-training,” *medRxiv*, pp. 2023–01, 2023.
- [6] C. Liu, S. Cheng, C. Chen, M. Qiao, W. Zhang, A. Shah, W. Bai, and R. Arcucci, “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*. Springer, 2023, pp. 637–647.
- [7] E. Tiu, E. Talius, P. Patel, C. P. Langlotz, A. Y. Ng, and P. Rajpurkar, “Expert-level detection of pathologies from unannotated chest x-ray images via self-supervised learning,” *Nature Biomedical Engineering*, pp. 1–8, 2022.
- [8] S.-C. Huang, L. Shen, M. P. Lungren, and S. Yeung, “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*, 2021, pp. 3942–3951.
- [9] F. Wang, Y. Zhou, S. Wang, V. Vardhanabhuti, and L. Yu, “Multi-granularity cross-modal alignment for generalized medical visual representation learning,” *arXiv preprint arXiv:2210.06044*, 2022.
- [10] C. Liu, A. Shah, W. Bai, and R. Arcucci, “Utilizing synthetic data for medical vision-language pre-training: Bypassing the need for real images,” *arXiv preprint arXiv:2310.07027*, 2023.
- [11] P. Cheng, L. Lin, J. Lyu, Y. Huang, W. Luo, and X. Tang, “Prior: Prototype representation joint learning from medical images and reports,” *arXiv preprint arXiv:2307.12577*, 2023.
- [12] H.-Y. Zhou, C. Lian, L. Wang, and Y. Yu, “Advancing radiograph representation learning with masked record modeling,” in *The Eleventh International Conference on Learning Representations*.
- [13] Y. Liu, S. Zhang, J. Chen, K. Chen, and D. Lin, “Pixmim: Rethinking pixel reconstruction in masked image modeling,” *arXiv preprint arXiv:2303.02416*, 2023.
- [14] K. He, X. Chen, S. Xie, Y. Li, P. Dollár, and R. Girshick, “Masked autoencoders are scalable vision learners,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2022, pp. 16 000–16 009.

- [15] Y. Liu, S. Zhang, J. Chen, Z. Yu, K. Chen, and D. Lin, “Improving pixel-based mim by reducing wasted modeling capability,” in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2023, pp. 5361–5372.
- [16] L. H. Li, P. Zhang, H. Zhang, J. Yang, C. Li, Y. Zhong, L. Wang, L. Yuan, L. Zhang, J.-N. Hwang *et al.*, “Grounded language-image pre-training,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2022, pp. 10965–10975.
- [17] Y. Gao, J. Liu, Z. Xu, J. Zhang, K. Li, R. Ji, and C. Shen, “Pyramidclip: Hierarchical feature alignment for vision-language model pretraining,” *Advances in neural information processing systems*, vol. 35, pp. 35959–35970, 2022.
- [18] C. Zhou, C. C. Loy, and B. Dai, “Extract free dense labels from clip,” in *European Conference on Computer Vision*. Springer, 2022, pp. 696–712.
- [19] H. Luo, J. Bao, Y. Wu, X. He, and T. Li, “Segclip: Patch aggregation with learnable centers for open-vocabulary semantic segmentation,” in *International Conference on Machine Learning*. PMLR, 2023, pp. 23033–23044.
- [20] Z. Wan, C. Liu, M. Zhang, J. Fu, B. Wang, S. Cheng, L. Ma, C. Quilodrán-Casas, and R. Arcucci, “Med-unic: Unifying cross-lingual medical vision-language pre-training by diminishing bias,” *arXiv preprint arXiv:2305.19894*, 2023.
- [21] X. Zhang, C. Wu, Y. Zhang, W. Xie, and Y. Wang, “Knowledge-enhanced visual-language pre-training on chest radiology images,” *Nature Communications*, vol. 14, no. 1, p. 4542, 2023.
- [22] W. Huang, H. Zhou, C. Li, H. Yang, J. Liu, and S. Wang, “Enhancing representation in radiography-reports foundation model: A granular alignment algorithm using masked contrastive learning,” *arXiv preprint arXiv:2309.05904*, 2023.
- [23] J. Ma, Y. He, F. Li, L. Han, C. You, and B. Wang, “Segment anything in medical images,” *Nature Communications*, vol. 15, no. 1, p. 654, 2024.
- [24] A. E. Johnson, T. J. Pollard, N. R. Greenbaum, M. P. Lungren, C.-y. Deng, Y. Peng, Z. Lu, R. G. Mark, S. J. Berkowitz, and S. Horng, “Mimic-cxr-jpg, a large publicly available database of labeled chest radiographs,” *arXiv preprint arXiv:1901.07042*, 2019.
- [25] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, Ł. Kaiser, and I. Polosukhin, “Attention is all you need,” *Advances in neural information processing systems*, vol. 30, 2017.
- [26] C. Ouyang, C. Biffi, C. Chen, T. Kart, H. Qiu, and D. Rueckert, “Self-supervision with superpixels: Training few-shot medical image segmentation without annotation,” in *Computer Vision—ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XXIX 16*. Springer, 2020, pp. 762–780.
- [27] F. Isensee, P. F. Jaeger, S. A. Kohl, J. Petersen, and K. H. Maier-Hein, “nnu-net: a self-configuring method for deep learning-based biomedical image segmentation,” *Nature methods*, vol. 18, no. 2, pp. 203–211, 2021.
- [28] C. Tomasi and R. Manduchi, “Bilateral filtering for gray and color images,” in *Sixth international conference on computer vision (IEEE Cat. No. 98CH36271)*. IEEE, 1998, pp. 839–846.
- [29] A. E. Johnson, T. J. Pollard, N. R. Greenbaum, M. P. Lungren, C.-y. Deng, Y. Peng, Z. Lu, R. G. Mark, S. J. Berkowitz, and S. Horng, “Mimic-cxr-jpg, a large publicly available database of labeled chest radiographs,” *arXiv preprint arXiv:1901.07042*, 2019.
- [30] E. Alsentzer, J. R. Murphy, W. Boag, W.-H. Weng, D. Jin, T. Naumann, and M. McDermott, “Publicly available clinical bert embeddings,” *arXiv preprint arXiv:1904.03323*, 2019.
- [31] G. Shih, C. C. Wu, S. S. Halabi, M. D. Kohli, L. M. Prevedello, T. S. Cook, A. Sharma, J. K. Amorosa, V. Arteaga, M. Galperin-Aizenberg *et al.*, “Augmenting the national institutes of health chest radiograph dataset with expert annotations of possible pneumonia,” *Radiology: Artificial Intelligence*, vol. 1, no. 1, p. e180041, 2019.
- [32] C. Steven G. Langer, PhD and M. George Shih, MD, “Siim-acr pneumothorax segmentation,” 2019.
- [33] J. Healthcare, “Object-cxr-automatic detection of foreign objects on chest x-rays,” 2020.
- [34] J. Redmon and A. Farhadi, “Yolov3: An incremental improvement,” *arXiv preprint arXiv:1804.02767*, 2018.
- [35] J. Irvin, P. Rajpurkar, M. Ko, Y. Yu, S. Ciurea-Ilcus, C. Chute, H. Marklund, B. Haghgoo, R. Ball, K. Shpanskaya *et al.*, “Chexpert: A large chest radiograph dataset with uncertainty labels and expert comparison,” in *Proceedings of the AAAI conference on artificial intelligence*, vol. 33, 2019, pp. 590–597.

- [36] X. Wang, Y. Peng, L. Lu, Z. Lu, M. Bagheri, and R. M. Summers, "Chestx-ray8: Hospital-scale chest x-ray database and benchmarks on weakly-supervised classification and localization of common thorax diseases," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2017, pp. 2097–2106.
- [37] L. Wang, Z. Q. Lin, and A. Wong, "Covid-net: A tailored deep convolutional neural network design for detection of covid-19 cases from chest x-ray images," *Scientific reports*, vol. 10, no. 1, pp. 1–12, 2020.
- [38] H. Zhang, P. Zhang, X. Hu, Y.-C. Chen, L. Li, X. Dai, L. Wang, L. Yuan, J.-N. Hwang, and J. Gao, "Glipv2: Unifying localization and vision-language understanding," *Advances in Neural Information Processing Systems*, vol. 35, pp. 36 067–36 080, 2022.
- [39] T.-Y. Lin, P. Dollár, R. Girshick, K. He, B. Hariharan, and S. Belongie, "Feature pyramid networks for object detection," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2017, pp. 2117–2125.
- [40] B. Boecking, N. Usuyama, S. Bannur, D. C. Castro, A. Schwaighofer, S. Hyland, M. Wetscherek, T. Naumann, A. Nori, J. Alvarez-Valle *et al.*, "Making the most of text semantics to improve biomedical vision-language processing," in *European conference on computer vision*. Springer, 2022, pp. 1–21.
- [41] C. Wu, X. Zhang, Y. Zhang, Y. Wang, and W. Xie, "Medklip: Medical knowledge enhanced language-image pre-training," *medRxiv*, pp. 2023–01, 2023.
- [42] M. Dombrowski, H. Reynaud, M. Baugh, and B. Kainz, "Foreground-background separation through concept distillation from generative image foundation models," in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2023, pp. 988–998.
- [43] A. Saporta, X. Gui, A. Agrawal, A. Pareek, S. Q. Truong, C. D. Nguyen, V.-D. Ngo, J. Seekins, F. G. Blankenberg, A. Y. Ng *et al.*, "Benchmarking saliency methods for chest x-ray interpretation," *Nature Machine Intelligence*, vol. 4, no. 10, pp. 867–878, 2022.

A Appendix / supplemental material

A.1 Limitations and Future Work

Our work primarily concentrates on learning dense visual representations from pseudo masks, which are generated from attention masks under language supervision. Due to the weak supervision signal, the pseudo masks may not effectively associate each pixel with the corresponding text tokens, potentially capping the performance of our method. Currently, our approach involves learning both global and pixel-level representations through VLP. In future studies, we aim to delve into regional visual representations during VLP to establish more precise correlations between specific chest X-ray (CXR) regions and phrases in radiology reports.

A.2 Broader Impacts

Our G2D model offers an effective approach for the automatic diagnosis of chest X-ray abnormalities using a small amount of annotated data. This can help decrease the burden on radiologists and enhance healthcare in underprivileged regions. However, medical data, such as chest X-rays and radiology reports, might include sensitive or potentially harmful information. We strongly advise a thorough examination of the data prior to using our model in real-world applications.

A.3 Pre-training Implementation Details




Figure 3: An exemplar pair of X-ray image and associated clinical report from the MIMIC-CXR dataset [24].

The chest X-ray (CXR) images from the MIMIC-CXR dataset [29] are resized to dimensions of 256×256 and subsequently center-cropped to 224×224 , adhering to the procedure described in [4, 8, 9], with an example shown in Fig 3. The intensity of each image is normalized to a range of $[0, 1]$. During the pre-training stage, we employ data augmentation techniques including random grayscale, random perspective, and auto contrast adjustments, using the PyTorch vision library¹.

A.4 Downstream Task Implementation Details

The data split information into train/valid/test sets are described in Tab. 6. For all downstream tasks, except from zero-shot image classification and visual grounding, we train with 1%, 10%, 100% of the training set. The downstream tasks are deployed on a 40G A100 GPU.

¹<https://pytorch.org/vision/stable/transforms.html>

Table 6: Details on Data Split: The symbol ‘/’ denotes that training/validation data is not required for the zero-shot tasks.

Task	Dataset	Split	Train	Valid	Test
Linear Classification	CheXpert [35]	[35]	186,027	5,000	202
	RSNA [31]	[9, 31]	16,010	5,337	5,337
	COVIDx [37]	[9, 37]	23988	5998	400
Fine-tuned Classification	CXR14 [36]	[21]	77,872	8,652	25,596
Image Segmentation	RSNA [31]	[8, 9]	16,010	5,337	5,337
	SIIM [32]	[8, 9]	8,433	1,807	1,807
Object Detection	RSNA [31]	[8, 9]	16,010	5,337	5,337
	Object-CXR [33]	[9]	6,400	1,600	1,000
Zero-shot Image Classification	RSNA [31]	[5]	/	/	5,337
	SIIM [32]	[5]	/	/	1,807
	CXR14 [36]	[36, 21]	/	/	25,596
	CheXpert [35]	[43, 21]	/	/	500
Zero-shot Visual Grounding	RSNA [31]	[5]	/	/	5,337
	SIIM [32]	[5]	/	/	1,807

A.4.1 Visual Localization

Medical Image Segmentation. For the segmentation tasks on the RSNA [31] and SIIM [32] datasets, we initially employ the vision encoder from the pre-trained model. Additionally, we transfer both the vision encoder and decoder from the pre-trained model, and proceed to train the segmentation network. We implement early stopping during the training process, limiting it to 50 epochs. A learning rate of $2e-4$ and a weight decay of 0.05 are adopted. AdamW is utilized as the optimizer, with β_1 and β_2 values set at 0.9 and 0.999, respectively. For the SIIM [32] dataset, the default batch size is set at 8, while for the RSNA [31] dataset, it is set at 16. All configurations strictly adhere to the protocol provided in [9].

Medical Image Object Detection. The pneumonia detection task on RSNA [31] and foreign objects detection task on Object-CXR [33] datasets are executed on a single A100 GPU. For both datasets, early stopping is implemented during the training process, limited to 50 epochs. AdamW is employed as the optimizer across both datasets. For the RSNA [31] dataset, a batch size of 8 is set for 1% data fine-tuning with a learning rate of $2e-4$, a weight decay of $1e-6$, and β_1, β_2 values of 0.9 and 0.999, respectively. For 10% and 100% data fine-tuning, the batch size is adjusted to 16, with a learning rate of $5e-4$, a weight decay of $1e-6$, and the same β_1, β_2 values. Similarly, for the Object-CXR [33] dataset, a batch size of 8 is set for 1% data fine-tuning, with the identical learning rate, weight decay, and β values as the RSNA dataset. For 10% and 100% data fine-tuning, the batch size is adjusted to 16, again with a learning rate of $5e-4$, a weight decay of $1e-6$, and β_1, β_2 values of 0.9 and 0.999. The IOU and NMS thresholds are set at [0.4, 0.45, 0.5, 0.55, 0.6, 0.65, 0.7, 0.75] and 0.5, respectively. All configurations are in strict compliance with the protocol delineated in [9].

A.4.2 Vision-Language Understanding

Zero-shot Image Classification. The original CXR image goes through a two-step preprocessing routine. Initially, it is resized to the dimension of 256×256 , and then center cropped to 224×224 . Following the methodologies outlined in [8, 9], all pixel values are normalized to the range [0, 1]. The resulting resized image is then fed through a visual encoder, followed by a visual projector to generate the image embedding \hat{v}_i . Simultaneously, the prompts are fed into a text encoder to obtain text embeddings \hat{i}_i . The classification evaluation hinges on measuring the cosine similarity between the image and text embeddings for each prompt associated with a specific class. The classification outcome is determined by comparing the cosine similarities. Specifically, if the cosine similarity between the image embedding and the positive prompt (e.g., *disease*) surpasses that between the image embedding and the corresponding negative prompt (e.g., *No disease*), the outcome is deemed positive. Conversely, if the reverse holds true, the outcome is negative. The prompt is designed following [7].

Zero-shot Visual Grounding. To execute this task, we adhere to the BioViL pipeline as described in [40]. The visual grounding task can be regarded as a pixel-level classification task, driven by the

text prompt and the dense visual embedding. The image is fed into the visual encoder to acquire the dense feature map \mathbf{V}_i from the final convolutional layer of the image encoder, yielding a shape of $C \times H \times W$. At the same time, the prompt is processed through the text encoder and projected into the cross-modal space, resulting in $\hat{\mathbf{l}}_i$. The cosine similarity between $\hat{\mathbf{l}}_i$ and all elements of \mathbf{V}_i at the channel level generates a similarity map. This map is then resized to match the original image size and utilized as the segmentation results to evaluate the zero-shot grounding performance.

A.4.3 Visual Recognition

We conduct evaluations on the CheXpert [35], RSNA [31], COVIDx [37], and CXR14 datasets [36]. In alignment with previous studies [8, 4, 9, 5, 21], linear classification is implemented on CheXpert [35], RSNA [31], and COVIDx [37]. Here, we update a randomly initialized linear layer while keeping the visual encoder frozen. We adhere to the official test set partition from [5, 21, 36] for a fair comparison. During our linear classification task, training is performed over 50 epochs with a learning rate of $5e-4$, a batch size of 8, employing the AdamW optimizer with parameters: $\beta_1 = 0.9$ and $\beta_2 = 0.999$. For the CXR14 dataset [36], we follow the experimental setup from [21], employing fine-tuned classification while updating all parameters from the visual encoder and linear layer. Images are resized to 256×256 and data augmentation is carried out as recommended in [21]. The AdamW optimizer is utilized with a learning rate of 1×10^{-4} and a batch size of 64 for 50 epochs. The linear classification tasks are executed on a single A100 GPU with 40GB memory, using the vision encoder from our pre-trained model as the visual backbone. Fine-tuning is carried out on the randomly initialized linear layer for 50 epochs with early stopping, maintaining a learning rate of $5e-4$ and a default batch size of 8. We set AdamW as our optimizer, with β_1 of 0.9, β_2 of 0.999, and a weight decay rate of $1e-6$.

A.5 Comparison under MedKLIP Configuration

Table 7: Performance of CXR14 Classification Fine-Tuning and Segmentation Results on SIIM and RSNA using the MedKLIP Setting [5].

	CXR14 (AUC)			RSNA (Dice)			SIIM (Dice)		
	1%	10%	100%	1%	10%	100%	1%	10%	100%
MedKLIP*	77.2	78.9	83.2	70.6	71.6	75.8	66.6	72.1	79.4
G2D(Ours)	80.4	83.8	86.1	73.8	76.1	76.5	70.6	74.5	82.3

To strictly compare our work with MedKLIP [5], we reimplement G2D for fine-tuning on CXR14 classification, as well as SIIM and RSNA segmentation tasks, adhering strictly to the MedKLIP configuration. This approach is necessary because the settings of MedKLIP differ significantly from the other methods that we compare to, such as [4, 8, 9, 6, 21]. Specifically, MedKLIP updates both the encoder and decoder during segmentation tasks, whereas the other methods only update the decoder and keep the encoder frozen. Moreover, MedKLIP employs its own customized data split for CXR14 classification, contrasting with KAD [21], which uses the official CXR14 dataset split.

Given these differences, comparing other methods directly under the MedKLIP setting could be seen as unfair. Therefore, we conducted a separate comparison between G2D and MedKLIP using the MedKLIP setting. The results, presented in Tab 7, demonstrate that G2D outperforms MedKLIP across all tasks and data ratios, even within the MedKLIP setting.

A.6 Verifying Pseudo Segmentation with Semantic Meaning

To investigate whether the improvements in G2D come from learning dense visual features through pseudo segmentation (PS) or from treating PS as a regularization term during pre-training, we perturbed the semantic integrity of pseudo masks by randomly shuffling them on a sample-wise basis (*i.e.*, making images and pseudo masks unpaired). This operation detaches pseudo masks’ semantic connection to the original images, ensuring that the PS task does not learn correct semantic information, but still provide regularisation to the segmentation as the pseudo mask is relatively smooth. The results are presented in Table 8. G2D with uncorrupted pseudo masks in PS (ours) significantly outperforms the results from the shuffled alternative (unpaired images and pseudo masks), not only in

Mask Construction	SIIM Dice	RSNA mAP	CXR14 AUC
Pseudo Mask without Semantic Meaning (shuffled)	50.9±2.4	7.6±1.2	63.7±2.1
Pseudo Mask with Semantic Meaning (Ours)	65.6±1.7	15.9±0.8	79.1±1.2

Table 8: Perturbation on Pseudo Masks.

visual localisation task but also in visual recognition task. The improved performance demonstrate that the proposed G2D indeed learns transferable visual features thanks to the semantic information provided by the pseudo masks, rather than merely treating PS as a regularization mechanism.

A.7 Pseudo Mask Visualization

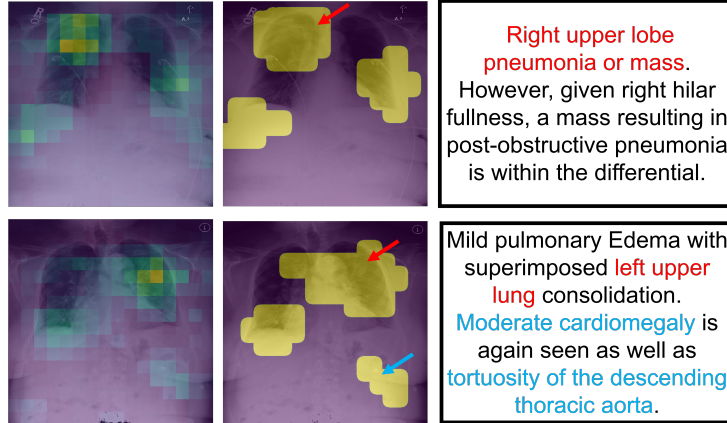


Figure 4: Pseudo Mask Visualization. **Left:** Aggregated attention map. **Middle:** Constructed pseudo mask for the pseudo segmentation task. **Red** and **blue** arrows point to areas related to specific text descriptions. **Right:** Corresponding radiology report. Red and blue text emphasize regions represented in the pseudo mask.

We visualize the aggregated attention map, pseudo mask, and paired medical reports in Fig 4. Intriguingly, without human annotations, both the attention map and pseudo mask successfully capture image regions corresponding to various report words. The pseudo masks manage to capture important parts of the image regions related to the highlighted words in the clinical reports, as indicated by the red and blue arrows in Fig 4. This suggests that the supervision signal for the PS pretext task is enriched by the clinical knowledge and high-level semantics, which explain why the PS pretext task may be better than the pixel reconstruction pretext task.

NeurIPS Paper Checklist

1. Claims

Question: Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope?

Answer: [Yes]

Justification: The main claims presented in the abstract and introduction accurately represent the contributions and scope of the paper.

Guidelines:

- The answer NA means that the abstract and introduction do not include the claims made in the paper.
- The abstract and/or introduction should clearly state the claims made, including the contributions made in the paper and important assumptions and limitations. A No or NA answer to this question will not be perceived well by the reviewers.
- The claims made should match theoretical and experimental results, and reflect how much the results can be expected to generalize to other settings.
- It is fine to include aspirational goals as motivation as long as it is clear that these goals are not attained by the paper.

2. Limitations

Question: Does the paper discuss the limitations of the work performed by the authors?

Answer: [Yes]

Justification: Refer to Section A.1.

Guidelines:

- The answer NA means that the paper has no limitation while the answer No means that the paper has limitations, but those are not discussed in the paper.
- The authors are encouraged to create a separate "Limitations" section in their paper.
- The paper should point out any strong assumptions and how robust the results are to violations of these assumptions (e.g., independence assumptions, noiseless settings, model well-specification, asymptotic approximations only holding locally). The authors should reflect on how these assumptions might be violated in practice and what the implications would be.
- The authors should reflect on the scope of the claims made, e.g., if the approach was only tested on a few datasets or with a few runs. In general, empirical results often depend on implicit assumptions, which should be articulated.
- The authors should reflect on the factors that influence the performance of the approach. For example, a facial recognition algorithm may perform poorly when image resolution is low or images are taken in low lighting. Or a speech-to-text system might not be used reliably to provide closed captions for online lectures because it fails to handle technical jargon.
- The authors should discuss the computational efficiency of the proposed algorithms and how they scale with dataset size.
- If applicable, the authors should discuss possible limitations of their approach to address problems of privacy and fairness.
- While the authors might fear that complete honesty about limitations might be used by reviewers as grounds for rejection, a worse outcome might be that reviewers discover limitations that aren't acknowledged in the paper. The authors should use their best judgment and recognize that individual actions in favor of transparency play an important role in developing norms that preserve the integrity of the community. Reviewers will be specifically instructed to not penalize honesty concerning limitations.

3. Theory Assumptions and Proofs

Question: For each theoretical result, does the paper provide the full set of assumptions and a complete (and correct) proof?

Answer: [NA]

Justification: This work mainly includes empirical contributions.

Guidelines:

- The answer NA means that the paper does not include theoretical results.
- All the theorems, formulas, and proofs in the paper should be numbered and cross-referenced.
- All assumptions should be clearly stated or referenced in the statement of any theorems.
- The proofs can either appear in the main paper or the supplemental material, but if they appear in the supplemental material, the authors are encouraged to provide a short proof sketch to provide intuition.
- Inversely, any informal proof provided in the core of the paper should be complemented by formal proofs provided in appendix or supplemental material.
- Theorems and Lemmas that the proof relies upon should be properly referenced.

4. Experimental Result Reproducibility

Question: Does the paper fully disclose all the information needed to reproduce the main experimental results of the paper to the extent that it affects the main claims and/or conclusions of the paper (regardless of whether the code and data are provided or not)?

Answer: [\[Yes\]](#)

Justification: We provide detailed experimental configurations in Sections 4.1, A.3, and A.4. Our code will be released after acceptance.

Guidelines:

- The answer NA means that the paper does not include experiments.
- If the paper includes experiments, a No answer to this question will not be perceived well by the reviewers: Making the paper reproducible is important, regardless of whether the code and data are provided or not.
- If the contribution is a dataset and/or model, the authors should describe the steps taken to make their results reproducible or verifiable.
- Depending on the contribution, reproducibility can be accomplished in various ways. For example, if the contribution is a novel architecture, describing the architecture fully might suffice, or if the contribution is a specific model and empirical evaluation, it may be necessary to either make it possible for others to replicate the model with the same dataset, or provide access to the model. In general, releasing code and data is often one good way to accomplish this, but reproducibility can also be provided via detailed instructions for how to replicate the results, access to a hosted model (e.g., in the case of a large language model), releasing of a model checkpoint, or other means that are appropriate to the research performed.
- While NeurIPS does not require releasing code, the conference does require all submissions to provide some reasonable avenue for reproducibility, which may depend on the nature of the contribution. For example
 - (a) If the contribution is primarily a new algorithm, the paper should make it clear how to reproduce that algorithm.
 - (b) If the contribution is primarily a new model architecture, the paper should describe the architecture clearly and fully.
 - (c) If the contribution is a new model (e.g., a large language model), then there should either be a way to access this model for reproducing the results or a way to reproduce the model (e.g., with an open-source dataset or instructions for how to construct the dataset).
 - (d) We recognize that reproducibility may be tricky in some cases, in which case authors are welcome to describe the particular way they provide for reproducibility. In the case of closed-source models, it may be that access to the model is limited in some way (e.g., to registered users), but it should be possible for other researchers to have some path to reproducing or verifying the results.

5. Open access to data and code

Question: Does the paper provide open access to the data and code, with sufficient instructions to faithfully reproduce the main experimental results, as described in supplemental material?

Answer: [Yes]

Justification: Our experiments are all conducted on publicly accessible datasets, and all experiment details are illustrated in Sections 4.1, A.3, and A.4. For experiment implementation, we follow the official code of existing works, all code can be found in their official GitHub repository.

Guidelines:

- The answer NA means that paper does not include experiments requiring code.
- Please see the NeurIPS code and data submission guidelines (<https://nips.cc/public/guides/CodeSubmissionPolicy>) for more details.
- While we encourage the release of code and data, we understand that this might not be possible, so “No” is an acceptable answer. Papers cannot be rejected simply for not including code, unless this is central to the contribution (e.g., for a new open-source benchmark).
- The instructions should contain the exact command and environment needed to run to reproduce the results. See the NeurIPS code and data submission guidelines (<https://nips.cc/public/guides/CodeSubmissionPolicy>) for more details.
- The authors should provide instructions on data access and preparation, including how to access the raw data, preprocessed data, intermediate data, and generated data, etc.
- The authors should provide scripts to reproduce all experimental results for the new proposed method and baselines. If only a subset of experiments are reproducible, they should state which ones are omitted from the script and why.
- At submission time, to preserve anonymity, the authors should release anonymized versions (if applicable).
- Providing as much information as possible in supplemental material (appended to the paper) is recommended, but including URLs to data and code is permitted.

6. Experimental Setting/Details

Question: Does the paper specify all the training and test details (e.g., data splits, hyperparameters, how they were chosen, type of optimizer, etc.) necessary to understand the results?

Answer: [Yes]

Justification: All experiment details are illustrated in Sections 4.1, A.3, and A.4.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The experimental setting should be presented in the core of the paper to a level of detail that is necessary to appreciate the results and make sense of them.
- The full details can be provided either with the code, in appendix, or as supplemental material.

7. Experiment Statistical Significance

Question: Does the paper report error bars suitably and correctly defined or other appropriate information about the statistical significance of the experiments?

Answer: [Yes]

Justification: We report the error bars for all ablation studies.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The authors should answer "Yes" if the results are accompanied by error bars, confidence intervals, or statistical significance tests, at least for the experiments that support the main claims of the paper.

- The factors of variability that the error bars are capturing should be clearly stated (for example, train/test split, initialization, random drawing of some parameter, or overall run with given experimental conditions).
- The method for calculating the error bars should be explained (closed form formula, call to a library function, bootstrap, etc.)
- The assumptions made should be given (e.g., Normally distributed errors).
- It should be clear whether the error bar is the standard deviation or the standard error of the mean.
- It is OK to report 1-sigma error bars, but one should state it. The authors should preferably report a 2-sigma error bar than state that they have a 96% CI, if the hypothesis of Normality of errors is not verified.
- For asymmetric distributions, the authors should be careful not to show in tables or figures symmetric error bars that would yield results that are out of range (e.g. negative error rates).
- If error bars are reported in tables or plots, The authors should explain in the text how they were calculated and reference the corresponding figures or tables in the text.

8. Experiments Compute Resources

Question: For each experiment, does the paper provide sufficient information on the computer resources (type of compute workers, memory, time of execution) needed to reproduce the experiments?

Answer: [Yes]

Justification: Refer to the first part of Sections 4.1, A.3, and A.4.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The paper should indicate the type of compute workers CPU or GPU, internal cluster, or cloud provider, including relevant memory and storage.
- The paper should provide the amount of compute required for each of the individual experimental runs as well as estimate the total compute.
- The paper should disclose whether the full research project required more compute than the experiments reported in the paper (e.g., preliminary or failed experiments that didn't make it into the paper).

9. Code Of Ethics

Question: Does the research conducted in the paper conform, in every respect, with the NeurIPS Code of Ethics <https://neurips.cc/public/EthicsGuidelines>?

Answer: [Yes]

Justification: This work is conducted in accordance with the NeurIPS Code of Ethics.

Guidelines:

- The answer NA means that the authors have not reviewed the NeurIPS Code of Ethics.
- If the authors answer No, they should explain the special circumstances that require a deviation from the Code of Ethics.
- The authors should make sure to preserve anonymity (e.g., if there is a special consideration due to laws or regulations in their jurisdiction).

10. Broader Impacts

Question: Does the paper discuss both potential positive societal impacts and negative societal impacts of the work performed?

Answer: [Yes]

Justification: Please refer to the Section A.2.

Guidelines:

- The answer NA means that there is no societal impact of the work performed.
- If the authors answer NA or No, they should explain why their work has no societal impact or why the paper does not address societal impact.

- Examples of negative societal impacts include potential malicious or unintended uses (e.g., disinformation, generating fake profiles, surveillance), fairness considerations (e.g., deployment of technologies that could make decisions that unfairly impact specific groups), privacy considerations, and security considerations.
- The conference expects that many papers will be foundational research and not tied to particular applications, let alone deployments. However, if there is a direct path to any negative applications, the authors should point it out. For example, it is legitimate to point out that an improvement in the quality of generative models could be used to generate deepfakes for disinformation. On the other hand, it is not needed to point out that a generic algorithm for optimizing neural networks could enable people to train models that generate Deepfakes faster.
- The authors should consider possible harms that could arise when the technology is being used as intended and functioning correctly, harms that could arise when the technology is being used as intended but gives incorrect results, and harms following from (intentional or unintentional) misuse of the technology.
- If there are negative societal impacts, the authors could also discuss possible mitigation strategies (e.g., gated release of models, providing defenses in addition to attacks, mechanisms for monitoring misuse, mechanisms to monitor how a system learns from feedback over time, improving the efficiency and accessibility of ML).

11. Safeguards

Question: Does the paper describe safeguards that have been put in place for responsible release of data or models that have a high risk for misuse (e.g., pretrained language models, image generators, or scraped datasets)?

Answer: [NA]

Justification: This work does not focus on content generation and uses clinically verified datasets for all experiments.

Guidelines:

- The answer NA means that the paper poses no such risks.
- Released models that have a high risk for misuse or dual-use should be released with necessary safeguards to allow for controlled use of the model, for example by requiring that users adhere to usage guidelines or restrictions to access the model or implementing safety filters.
- Datasets that have been scraped from the Internet could pose safety risks. The authors should describe how they avoided releasing unsafe images.
- We recognize that providing effective safeguards is challenging, and many papers do not require this, but we encourage authors to take this into account and make a best faith effort.

12. Licenses for existing assets

Question: Are the creators or original owners of assets (e.g., code, data, models), used in the paper, properly credited and are the license and terms of use explicitly mentioned and properly respected?

Answer: [Yes]

Justification: Please refer to Section 4.1.

Guidelines:

- The answer NA means that the paper does not use existing assets.
- The authors should cite the original paper that produced the code package or dataset.
- The authors should state which version of the asset is used and, if possible, include a URL.
- The name of the license (e.g., CC-BY 4.0) should be included for each asset.
- For scraped data from a particular source (e.g., website), the copyright and terms of service of that source should be provided.

- If assets are released, the license, copyright information, and terms of use in the package should be provided. For popular datasets, paperswithcode.com/datasets has curated licenses for some datasets. Their licensing guide can help determine the license of a dataset.
- For existing datasets that are re-packaged, both the original license and the license of the derived asset (if it has changed) should be provided.
- If this information is not available online, the authors are encouraged to reach out to the asset's creators.

13. **New Assets**

Question: Are new assets introduced in the paper well documented and is the documentation provided alongside the assets?

Answer: [NA]

Justification: There is no new assets released in this work.

Guidelines:

- The answer NA means that the paper does not release new assets.
- Researchers should communicate the details of the dataset/code/model as part of their submissions via structured templates. This includes details about training, license, limitations, etc.
- The paper should discuss whether and how consent was obtained from people whose asset is used.
- At submission time, remember to anonymize your assets (if applicable). You can either create an anonymized URL or include an anonymized zip file.

14. **Crowdsourcing and Research with Human Subjects**

Question: For crowdsourcing experiments and research with human subjects, does the paper include the full text of instructions given to participants and screenshots, if applicable, as well as details about compensation (if any)?

Answer: [NA]

Justification: This work has no human subjects.

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Including this information in the supplemental material is fine, but if the main contribution of the paper involves human subjects, then as much detail as possible should be included in the main paper.
- According to the NeurIPS Code of Ethics, workers involved in data collection, curation, or other labor should be paid at least the minimum wage in the country of the data collector.

15. **Institutional Review Board (IRB) Approvals or Equivalent for Research with Human Subjects**

Question: Does the paper describe potential risks incurred by study participants, whether such risks were disclosed to the subjects, and whether Institutional Review Board (IRB) approvals (or an equivalent approval/review based on the requirements of your country or institution) were obtained?

Answer: [NA]

Justification: This work has no human subjects.

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Depending on the country in which research is conducted, IRB approval (or equivalent) may be required for any human subjects research. If you obtained IRB approval, you should clearly state this in the paper.

- We recognize that the procedures for this may vary significantly between institutions and locations, and we expect authors to adhere to the NeurIPS Code of Ethics and the guidelines for their institution.
- For initial submissions, do not include any information that would break anonymity (if applicable), such as the institution conducting the review.