

SERPENT-VLM : Self-Refining Radiology Report Generation Using Vision Language Models

Manav Nitin Kapadnis** Sohan Patnaik* Abhilash Nandy Sourjyadip Ray
Pawan Goyal Debdoot Sheet

iammanavk@gmail.com sohanpatnaik106@gmail.com
Indian Institute of Technology Kharagpur
India

Abstract

Radiology Report Generation (R2Gen) demonstrates how Multi-modal Large Language Models (MLLMs) can automate the creation of accurate and coherent radiological reports. Existing methods often *hallucinate* details in text-based reports that don't accurately reflect the image content. To mitigate this, we introduce a novel strategy, **SERPENT-VLM** (**S**elf **R**efining **R**adiology **R**e**P**ort **G**ENERa**T**ion using **V**ision **L**anguage **M**odels), which improves the R2Gen task by integrating a self-refining mechanism into the MLLM framework. We employ a unique *self-supervised loss* that leverages similarity between pooled image representations and the contextual representations of the generated radiological text, alongside the standard Causal Language Modeling objective, to refine image-text representations. This allows the model to scrutinize and align the generated text through dynamic interaction between a given image and the generated text, therefore reducing hallucination and continuously enhancing nuanced report generation. SERPENT-VLM outperforms existing baselines such as LLaVA-Med, BiomedGPT, etc., achieving SoTA performance on the IU X-ray and Radiology Objects in COntext (ROCO) datasets, and also proves to be robust against noisy images. A qualitative case study emphasizes the significant advancements towards more sophisticated MLLM frameworks for R2Gen, opening paths for further research into self-supervised refinement in the medical imaging domain.

1 Introduction

Radiology Report Generation (R2Gen) serves as a crucial link between medical imaging and natural language processing, to automate the interpretation of radiological images into comprehensive text reports. This task requires models to learn long-range dependencies effectively while generating

the report, a challenge that remains largely unmet in current systems. The primary goal of R2Gen is to generate accurate and comprehensive medical reports from radiological imagery, an essential step toward enhancing diagnostic accuracy and efficiency. Prevailing methods (Vinyals et al., 2015; Xu et al., 2015; Tang et al., 2023; You et al., 2016; Tang et al., 2021) in R2Gen often rely on (1) large datasets for pre-training to impart domain-specific knowledge, and (2) typically utilizing compute-intensive encoder-decoder architectures for fine-tuning. These approaches are fraught with drawbacks, such as omission of minor yet clinically significant details (Wang et al., 2022b; You et al., 2021; Wang et al., 2021) and the persistent issue of *hallucination* as seen in Fig. 1, where generated reports from LLaVA-Med and BiomedGPT wrongly include details not present in the images. Minimizing hallucinations in radiology report generation is crucial since these inaccuracies can lead to misdiagnoses, directly impacting patient treatment plans and outcomes. Moreover, reducing hallucinations ensures the reliability and trustworthiness of automated reports, which is vital for maintaining clinical credibility and facilitating effective patient care. Therefore, the limitations pertaining to existing approaches underscore the necessity for a more refined approach for accurate medical diagnosis, addressing the critical gaps in R2Gen.

In this paper, we introduce a streamlined pipeline, SERPENT-VLM, which begins by processing a given X-ray image by passing it through a visual encoder and mapping it to a vector representation in a high-dimensional space. This process facilitates a nuanced understanding of the medical imagery. The encoded image, alongside a report generation prompt, is then passed as inputs to a Large Language Model (LLM) for text generation. We employ a cross-entropy loss for the **causal language modeling objective** and **introduce a novel self-refining objective** that leverages the pooled

* Equal contribution.

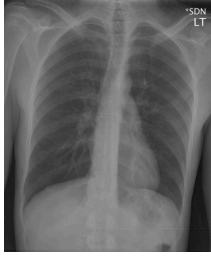
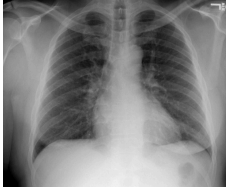
	Ground Truth	LlaVA-Med	BiomedGPT	SERPENT-VLM
	The lungs are hyperexpanded . cardiomeastinal silhouette is within normal limits . no pleural effusion focal airspace opacities or pneumothorax . no free subdiaphragmatic air .	clear and well-defined lung field. ribs are well-defined, no obvious fractures or dislocations. The costophrenic angles are sharp, and there are no pleural effusions. However, there appears to be a small, faint, and irregular shadow in the lower left lung field, which could be indicative of a small nodule or mass.	Lung fields are clear, with well-defined contours. Ribs show no fractures or dislocations, appearing intact. Costophrenic angles maintain sharpness, indicating absence of pleural effusions. A slight haziness is noted in the right upper lung field, potentially indicating early signs of infiltrative process.	heart size and mediastinal contours are within normal limits . pulmonary vascularity is within normal limits . no focal consolidation pleural effusion or pneumothorax identified . visualized osseous structures of the thorax are without acute abnormality .
	The cardiomeastinal silhouette and pulmonary vasculature are within normal limits in size. The lungs are clear of focal airspace disease, pneumothorax, or pleural effusion. There are no acute bony findings.	Clear, well-defined diaphragm. Slight irregularity, lower right lung lobe, potential small nodule, early pulmonary issue. Normal heart shadow, no cardiomegaly. Normal-sized aorta, no pneumothorax, no pleural effusion. Intact ribs, vertebral column, no fractures, dislocations. Centrally located mediastinum, no masses, lymphadenopathy. Normal lung vascular markings, no atelectasis, pneumonia.	clear, well-defined diaphragm and lung fields, except for faint linearity in the middle right lung, suggesting early fibrotic changes. The heart shadow is normal, excluding cardiomegaly, while the aorta is of normal size with no evidence of pneumothorax or pleural effusion. Bones, including ribs and vertebral column, show no fractures or dislocations, despite minimal displacement in the right clavicle.	the heart size and mediastinal contours are within normal limits . the lungs are clear without focal airspace opacity pleural effusion or pneumothorax . there is no visible pneumothorax . there is no evidence of acute bony abnormality

Figure 1: Generated report samples on IU-Xray dataset. We qualitatively analyze reports generated by medical pre-trained LLMs LlaVA-Med and BioMedGPT with SERPENT-VLM. Hallucinated information in the reports is highlighted using yellow.

image representation and the generated report’s contextual representation. This allows for tuning the network without compromising inference latency, while significantly improving performance evaluated using metrics such as *Bleu*, *Rouge_L*, *BertScore*.

The contributions of our work are summarized as follows:

1. Our approach **does not compromise on inference latency**, adopting a refining strategy through a novel loss function used only for fine-tuning
2. The introduction of a self-refining loss ensures the generation of nuanced, **hallucination-free** radiology reports
3. Our system not only matches but surpasses the performance of leading generalistic pre-trained medical LLMs.
4. Our approach demonstrates **robustness against noisy image** inputs, maintaining the generation of comprehensive reports.

This marks a substantial advancement in the field of R2Gen, setting new benchmarks for accuracy, efficiency, and robustness.

The remainder of the paper is organized as follows: We begin by delving into the literature review in Section 2, focusing on current and past state-of-the-art (SoTA) methodologies in the domain of radiological report generation. Section 3

discusses the proposed strategy for the self-refining fine-tuning our approach. The datasets, baselines, experimental setups, and ablation studies are detailed in Section 4. Finally, we conclude with a summary of our findings in Section 5.

2 Related Work

Medical Report Generation (MRG): Medical Report Generation has been extensively studied through ML models. (Jing et al., 2018) proposed a co-attention network that aligns visual and textual information to generate comprehensive radiology reports. Further enhancing the capabilities, a memory-driven transformer (Chen et al., 2020) integrates memory modules for encoding and decoding processes, allowing for more sophisticated report generation (Chen et al., 2020, 2021). Cross-modal learning (Wang et al., 2022a) utilizes prototype matrices and contrastive losses to refine the learning of visual-textual correlations, complemented by a self-boosting framework to align image features with report text (Wang et al., 2021). (Liu et al., 2021) addressed the problem of mitigating inherent biases through a data-driven method, introducing a prior-posterior knowledge-based report generation. (Nooralahzadeh et al., 2021) leveraged curriculum learning to extract global concepts to create a bridge between images and text. Task-specific architecture with sentence-level attention mechanism across visual features (Yuan et al., 2019) allows the model to capture key medical concepts from

images. A weakly supervised paradigm to amplify hard negative samples (Yan et al., 2021) addresses the medical data scarcity challenge.

Large Language Models and Vision language Models: The advent of Large Language Models (LLMs) such as GPT-4, Claude, BARD showcase excellent zero-shot language understanding (Bro, 2020; Li et al., 2021; Liu et al., 2021; Irvin et al., 2019); image understanding and visual question answering (Team et al., 2023) capabilities. Open-source LLMs, like LLaMA and BLOOM, and Multi-modal LLMs such as LLaVA (Liu et al., 2024), Open Flamingo (Awadalla et al., 2023) have also democratized access to cutting-edge generative technology (Ouyang et al., 2022; Pan et al., 2020). Furthermore, domain-specific models LLaVA-Med (Li et al., 2023) and BiomedGPT (Zhang et al., 2024) have shown promising results in pathology and radiology-related tasks. However, knowledge grounding for medical reports (Hyland et al., 2023), thereby reducing hallucination produced by these models remains a challenge.

Source & Representation of Feedback: Iterative refinement in MRG has traditionally relied on human feedback to achieve high-quality outputs (Tandon et al., 2022). Scalar reward functions and domain-specific feedback tools, such as compilers, were proposed as cost-effective alternatives to human feedback (Le et al., 2022; Yasunaga and Liang, 2020). Recent developments show that Large Language Models (LLMs) can self-evaluate their responses. However, applying this to Multi-modal Large Language Models remains largely unexplored in terms of generating grounded and hallucination-free responses.

We now discuss the proposed methodology in the subsequent section.

3 Methodology

3.1 Overview of SERPENT-VLM

We summarize the pipeline of SERPENT-VLM in Figure 2. It consists of two branches to establish the learning optimization criterion. **1) Causal Language Modeling Objective** enforces standard cross-entropy loss (step 4 in Fig. 2) for supervised radiology report generation. Our approach consists of a visual encoder that extracts information from chest X-ray images (step 1 in Fig. 2), a visual mapper that projects low dimensional image features onto high dimensional feature space (step 2 in Fig. 2) and a Large Language Model that au-

to regressively generates the diagnostic radiological report (step 3 in Fig. 2). To further reduce hallucination, we construct a pooled representation of the given X-ray image, a contextual representation leveraging the attention weights and last hidden states of the generated report and enforce **2) Self Refining Objective** that tries to maximise the similarity between pooled image representation and the contextual representation of the generated report through a self-supervised loss criterion (step 5 in Fig. 2). We train the network through a weighted combination of both the losses (step 6 in Fig. 2), thereby enabling SERPENT-VLM to continuously refine itself by aligning generated text with the input image. We now discuss the details of each component.

3.2 SERPENT-VLM Framework

The architecture of SERPENT-VLM can be partitioned into three different modules - a visual encoder, a visual mapper and a large language model (LLM). Formally, consider a chest X-ray image $I_v \in \mathbb{R}^{C \times H \times W}$, where C is the number of input channels, H , W being the height and width of the image respectively. $I_v = [I_{v_1}, I_{v_2}, \dots, I_{v_k}]$ comprises of a sequence of k patches with $I_{v_i} \in \mathbb{R}^{C \times P \times P}$ being the i^{th} patch, and P is the patch size. We leverage a transformer-based visual encoder V_{enc} to encode and obtain contextual representation $\tilde{e}_{v_i} \in \mathbb{R}^{d_v}$ denoted by Eq. 1 and aggregate each encoded patch to obtain a global image representation \tilde{e}_v depicted by Eq. 2.

$$\tilde{e}_{v_1}, \tilde{e}_{v_2}, \dots, \tilde{e}_{v_k} = V_{enc}(I_{v_1}, I_{v_2}, \dots, I_{v_k}) \quad (1)$$

$$\tilde{e}_v = V_{pooler}(\tilde{e}_{v_1}, \tilde{e}_{v_2}, \dots, \tilde{e}_{v_k}) \quad (2)$$

The encoded image features inherently reside in a visual feature space, which is distinct and not directly compatible with the textual feature space, and hence need to be aligned with the word embedding space of the LLM. To ensure this, we use a learnable visual mapper V_{map} to project the patch embeddings \tilde{e}_{v_i} onto the word embedding space. Formally, $e_{v_i} = V_{map}(\tilde{e}_{v_i})$. We construct a seed prompt T instructing the LLM to generate a report conditioned on the image I_v , and obtain the corresponding tokens $\mathcal{T}_{tokens} = [t_1, t_2, \dots, t_{|\mathcal{T}_{tokens}|}]$ which is given as input to the *Embedding* module of the LLM to construct the token embeddings (refer Eq. 3),

$$e_{t_1}, e_{t_2}, \dots, e_{t_{|\mathcal{T}_{tokens}|}} = Embedding(t_1, t_2, \dots, t_{|\mathcal{T}_{tokens}|}) \quad (3)$$

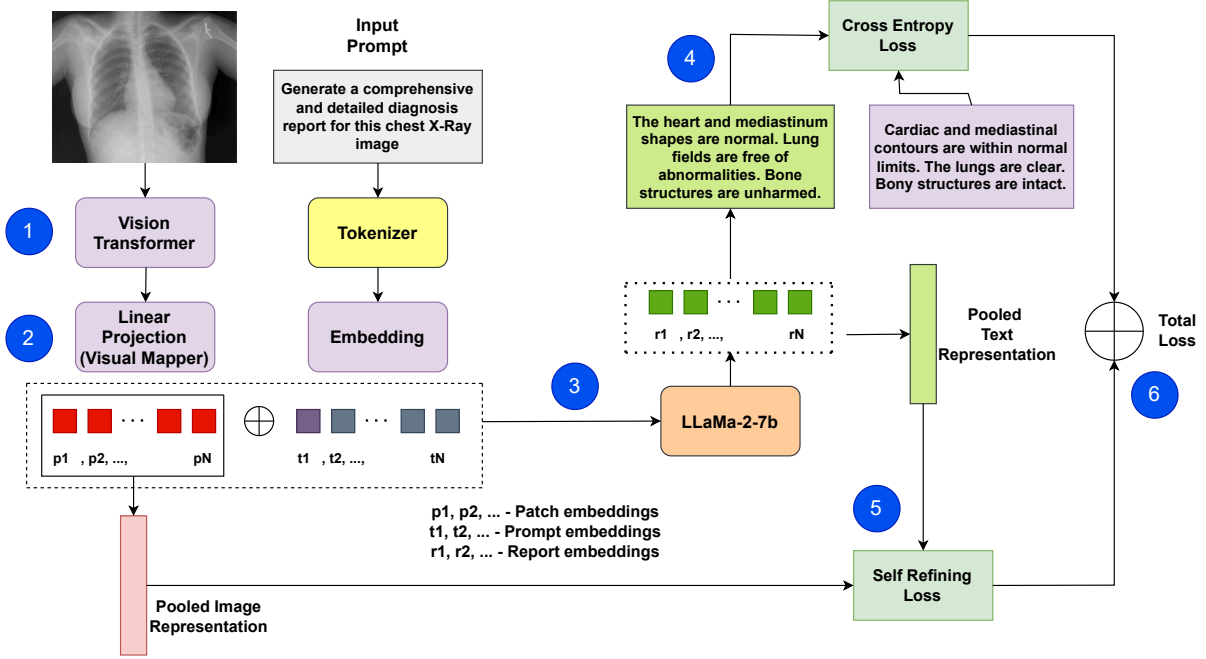


Figure 2: Overview of the SERPENT-VLM pipeline. The X-ray image is processed using a visual encoder (step 1) and projected onto a high-dimensional space using a visual mapper (step 2). The encoded image with the report generation prompt is fed into the LLM (step 3). Cross-entropy loss is employed (step 4) for the causal language modeling objective. The pooled image representation and the Contextual representation of the generated report are used to compute the self-refining loss (step 5). A weighted combination of both objectives is used to train the network (step 6).

We concatenate the sequence of projected image patch embeddings e_{v_i} with the seed prompt text embeddings e_{t_j} to obtain a sequence of input embeddings $e_{\mathcal{I}} = [e_v; e_t]$ which are given as input to the decoder-only LLM denoted by TD for generating the logits of the response tokens in autoregressive fashion. V_{enc} , V_{pooler} , V_{map} and TD are trained through cross-entropy loss \mathcal{L}_{report} enforced between the generated logits and the actual responses. To further guide the report generation process by aligning the generated response with the input image, we enforce a self-supervised *refining loss*.

3.3 Self-refining Strategy

We construct an aggregated representation of the generated text by utilizing the attention weights of the last layer of TD . Consider the logit distribution for each generated token as $l_i \in \mathbb{R}^d$, where d is the vocabulary size of TD . To encode the representation of each generated token, which is further used to compute the *self-refining loss* in a differentiable fashion, we leverage Gumbel-Softmax on the logit distribution to obtain \hat{l}_i for each predicted token. We construct the aggregated representation $\hat{e}_i^p = \sum_{j=1}^d e_j \hat{l}_{ij}$ of each predicted token by

taking a weighted sum of the embedding matrix $E = e_1, e_2, \dots, e_d$ with \hat{l}_i being the corresponding weights. Formally,

$$\hat{l}_{ij} = \frac{e^{(\log(l_{ij}) + g_{ij})/\tau}}{\sum_{j=1}^d e^{(\log(l_{ij}) + g_{ij})/\tau}} \quad (4)$$

Since, the gumbel-softmax operator makes the logit distribution peaky, taking a weighted sum effectively yields the predicted token embeddings. Further, we construct an aggregated representation $h_t \in \mathbb{R}^{d_t}$ of the predicted token embeddings by leveraging the attention weights from the last layer of TD . We hypothesize that aligning the aggregated representation of the generated report with the pooled input image representation would reduce hallucination and ground the report generation task. For this, we enforce a *self-refining loss* between h_t and e_v depicted by Eq. 3.3

$$\mathcal{L}_{refine} = \frac{1}{b} \sum_i^b e^{-h_t^T e_v}, \quad (5)$$

where b is the batch size.

Minimizing the negative exponential of the similarity between the image and generated text representation pushes the representation closer, thus further grounding the report generation process. We

optimize our network with a weighted combination of both the causal language modeling objective and the self-refining objective. The total loss is denoted by Eq. 6

$$\mathcal{L}_{total} = \lambda_{report} \mathcal{L}_{report} + \lambda_{refine} \mathcal{L}_{refine} \quad (6)$$

\mathcal{L}_{report} depicts the standard causal language modeling objective that ensures the conditional generation of radiological report text based on the input image, whereas \mathcal{L}_{refine} ensures that the generated report is grounded in context of the input image, thereby establishing a robust pipeline for radiology report generation.

4 Experiments and Evaluation

We now discuss the details corresponding to the experiments and ablation studies carried out and enumerate the observations.

4.1 Implementation Details

We discuss the technical details and hyperparameter settings for all the experiments. For the visual encoder V_{enc} , we employed the base version of Swin-Transformer-V2¹ and a feed-forward neural network for V_{map} . We leverage LLaMA2-7B² as our primary LLM. Further, the hidden dimension of d_v of V_{enc} and d_t of TD are 768 and 1024 respectively. We freeze the weights of V_{enc} , however keep V_{map} trainable. We employ LoRA with a rank and α -scaling factor of 16 each to fine-tune the underlying LLM TD . We train SERPENT-VLM for 15 epochs on IU-Xray dataset and 20 epochs on the ROCO dataset with mixed precision on an effective batch size (BS) of 6 using one NVIDIA A40 48GB GPU using a learning rate of 1×10^{-4} with linear rate scheduler through AdamW optimizer. For inference, we leverage beam search decoding with beam size configured to 3.

4.2 Datasets and Evaluation Metrics:

We evaluate SERPENT-VLM on two commonly used datasets diverse modality -

1. **IU X-Ray** which is a widely used publicly available dataset for medical report generation tasks containing 3,955 fully de-identified radiology reports with sections such as Impression, Findings, Indication, etc., each associated with frontal and/or lateral chest X-rays, totaling 7,470 images;

2. **ROCO** which has ‘radiology’ and ‘out-of-class’ subsets (synthetic radiology images, clinical photos, portraits, compound radiology images, and digital art) of roughly 65,460 and 8,182 ‘radiology’, and 4,902 and 613 ‘out-of-class’ images in the train and test set respectively.

Since the reports are verbose and need to be accurately measured with word-level precision, we compute overlap-based metrics like BLEU and Rouge-L, and a semantic similarity-based metric BertScore for evaluating the efficacy of our approach.

Dataset	Train	Val	Test	Image Views
IU X-Ray	2769	791	395	Frontal and Lateral
ROCO	65460	8183	8182	Frontal

Table 1: Statistics of Evaluation Datasets

4.3 Performance of SERPENT-VLM on Radiology Report Generation

Table 2 illustrates the comprehensive comparison of SERPENT-VLM against various state-of-the-art baselines across the IU-Xray and ROCO datasets. In comparison with traditional non-LLM approaches such as Show-Tell (Vinyals et al., 2015), Att2in (Xu et al., 2015), and R2Gen (Chen et al., 2020), SERPENT-VLM exhibits significant improvements. For instance, on the IU-Xray dataset, SERPENT-VLM achieves a $Bleu_4$ score of 0.190, surpassing Show-Tell’s 0.078 and R2Gen’s 0.165, and even outperforming the more advanced R2GenCMN, which scores 0.170. This indicates not only an improvement in capturing long-range dependencies but also a notable reduction in detail hallucination, a common issue in earlier models. Furthermore, when compared to Medical LLMs and generalistic Vision-Language Models such as LLaVA-Med (Li et al., 2023), BiomedGPT (Zhang et al., 2024), and MiniGPT4 (Zhu et al., 2023), SERPENT-VLM demonstrates superior performance, marking a significant leap in R2Gen. For example, against LLaVA-Med, which records a $Bleu_4$ of 0.186 on IU-Xray, SERPENT-VLM shows a marked improvement with a score of 0.190. Similarly, in the context of $BertScore$, SERPENT-VLM achieves an impressive 0.935 compared to LLaVA-Med’s 0.845 and BiomedGPT’s 0.793, underscoring its enhanced textual coherence.

¹<https://huggingface.co/microsoft/swinv2-base-patch4-window12-192-22k>

²<https://huggingface.co/meta-llama/Llama-2-7b-chat-hf>

Methods	IU-Xray						ROCO					
	Bleu ₁	Bleu ₂	Bleu ₃	Bleu ₄	Rouge _L	BertScore	Bleu ₁	Bleu ₂	Bleu ₃	Bleu ₄	Rouge _L	BertScore
Show-Tell	0.243	0.13	0.108	0.078	0.307	0.378	0.104	0.076	0.051	0.027	0.089	0.34
Att2in	0.248	0.134	0.116	0.091	0.309	0.386	0.106	0.077	0.052	0.027	0.091	0.347
AdaAtt	0.284	0.207	0.15	0.126	0.311	0.442	0.122	0.089	0.060	0.031	0.104	0.397
Transformer	0.372	0.251	0.147	0.136	0.317	0.579	0.159	0.116	0.079	0.041	0.137	0.521
M2transformer	0.402	0.284	0.168	0.143	0.328	0.626	0.172	0.125	0.085	0.044	0.148	0.563
R2Gen	0.47	0.304	0.219	0.165	0.371	0.732	0.201	0.147	0.099	0.052	0.173	0.658
R2GenCMN	0.475	0.309	0.222	0.17	0.375	0.74	0.169	0.148	0.100	0.052	0.175	0.665
MSAT	0.481	0.316	0.226	0.171	0.372	0.749	0.212	0.150	0.102	0.053	0.177	0.673
METransformer	0.483	0.322	0.228	0.172	0.38	0.752	0.211	0.151	0.102	0.053	0.178	0.676
R2GenGPT (Deep)	0.480	0.316	0.216	0.169	0.377	0.748	0.213	0.150	0.101	0.053	0.177	0.672
MiniGPT4	0.494	0.329	0.220	0.179	0.390	0.767	0.219	0.156	0.103	0.056	0.183	0.689
BiomedGPT	0.516	0.343	0.233	0.183	0.403	0.793	0.229	0.163	0.109	0.058	0.189	0.712
LlaVA-Med	0.528	0.346	0.237	0.186	0.422	0.845	0.234	0.164	0.111	0.061	0.198	0.759
SERPENT-VLM	0.547	0.356	0.242	0.190	0.452	0.935	0.243	0.169	0.108	0.057	0.212	0.84

Table 2: Results of SERPENT-VLM on Benchmark datasets

4.4 Discussion on the Impact of different Design Choices for SERPENT-VLM

We carry experiments pertaining to two different design choices for SERPENT-VLM and establish the efficacy of the proposed architecture through the comparative analysis across experiments.

1. Effect of relative importance of two losses:

We vary the relative importance self-refining loss (λ_{refine}) and report-generation loss (λ_{report}) in Eq. 6. Table 3 shows that combining the two losses yields much better performance for IU X-ray and ROCO compared to just using the report generation loss (row 5 vs. row 2). This highlights that self-refining loss complements the report generation loss by grounding the generated report on the input image, thereby reducing hallucination. Further, it is observed that using only self-refining loss (row 1) leads to a degradation in performance because SERPENT-VLM is trained only through a self-supervised paradigm without any kind of supervision. As observed, this equilibrium is not merely about avoiding hallucinations but also about fostering a synergistic effect where each loss component reinforces the other, thereby elevating the overall quality and reliability of the automated radiology reports. The findings from our experiments provide compelling evidence for the critical role of balanced loss parameters in achieving the desired outcomes, advocating for a nuanced approach in their application within the framework of SERPENT-VLM.

2. Effect of contextual representation design strategy:

We explore different aggregation

strategies for obtaining the contextual representation of the generated report. As depicted in Table 4, attention-based aggregation outperforms other aggregation strategies by a significant margin by obtaining a BertScore of 0.935 and 0.840; BLEU₁ score of 0.547 and 0.243 on IU X-ray and ROCO respectively. Average pooling (average of token representations), Max pooling (token representation with maximum L2-norm) and Top-k average pooling (average top $k = 5$ token representations based on attention-weights) give sub-optimal performance on both IU X-ray and ROCO benchmark, thereby establishing the critical importance of sophisticated feature integration methods in enhancing the model’s capability to synthesize coherent and contextually relevant radiology reports. Exploration into different aggregation strategies reveals that the sophistication and adaptability of the aggregation mechanism play a pivotal role in the efficacy of medical report generation models.

4.5 How robust is SERPENT-VLM to noisy images?

We assess the robustness of SoTA methods LlaVA-Med and BiomedGPT, with our method SERPENT-VLM, by introducing Gaussian noise to radiological images. Fig. 3 demonstrate that SERPENT-VLM significantly outperforms the current SoTA models, LlaVA-Med and BiomedGPT, across all Gaussian Noise scales, maintaining higher BLEU₁ (5-6% higher) and BertScore (9-10% higher) metrics, thus showcasing superior robustness in report generation under noisy and corrupted images. This also highlights SERPENT-VLM’s ability to focus

Dataset	λ_{Report}	λ_{Refine}	$Bleu_1$	$Bleu_2$	$Bleu_3$	$Bleu_4$	$Rouge_L$	$BertScore$
IU-Xray	0	1.0	0.416	0.270	0.184	0.144	0.344	0.711
	0.3	0.7	0.547	0.356	0.242	0.190	0.452	0.935
	0.5	0.5	0.492	0.320	0.218	0.171	0.407	0.842
	0.7	0.3	0.479	0.311	0.212	0.166	0.396	0.818
	1	0.0	0.451	0.311	0.200	0.157	0.373	0.771
ROCO	0	1	0.187	0.130	0.083	0.044	0.163	0.647
	0.3	0.7	0.243	0.169	0.108	0.057	0.212	0.840
	0.5	0.5	0.214	0.149	0.095	0.050	0.187	0.739
	0.7	0.3	0.207	0.144	0.092	0.048	0.180	0.714
	1	0	0.194	0.135	0.086	0.046	0.170	0.672

Table 3: Impact of combining self-refining loss (weight λ_{refine}) with report-generation loss (weight λ_{report}). Fusing both the loss components gives optimal performance.

Dataset	Design Strategy	$Bleu_1$	$Bleu_2$	$Bleu_3$	$Bleu_4$	$Rouge_L$	$BertScore$
IU-Xray	Attention based aggregation	0.547	0.356	0.242	0.190	0.452	0.935
	Average pooling	0.410	0.267	0.182	0.143	0.339	0.701
	Top k average pooling	0.465	0.303	0.206	0.162	0.384	0.795
	Max pooling	0.383	0.249	0.169	0.133	0.316	0.655
ROCO	Attention based aggregation	0.243	0.169	0.108	0.057	0.212	0.840
	Average pooling	0.190	0.132	0.084	0.044	0.165	0.655
	Top k average pooling	0.199	0.139	0.089	0.047	0.174	0.689
	Max pooling	0.170	0.118	0.076	0.040	0.148	0.588

Table 4: Performance comparison of different design strategies for contextual representation. Attention weights-based aggregation displays superior performance.

on relevant parts of the image, thereby mitigating the effects of added noise and grounding the generated report - an indication of reduction in hallucination phenomena. The integration of SERPENT-VLM could markedly enhance diagnostic accuracy, aiding radiologists in delivering faster and more accurate patient care.

5 Summary and Conclusion

In this paper, we propose SERPENT-VLM, an innovative method for producing detailed and accurate radiology reports from Chest X-rays without hallucinations. The process utilizes a frozen visual encoder to transform X-ray images into a high-dimensional space, which a Large Language Model (LLM) then uses to generate initial reports. These reports undergo further refinement through a novel combination of self-refining loss and Causal Language Modeling Loss, significantly surpassing existing methods as detailed in Section 4. Our experiments in Section 4 and supplementary materials, confirm the effectiveness of our self-refining approach, even with distorted noisy images. Our future works involve the extension of our method to other medical imaging types, such as MRIs and CT


scans, and to incorporate diagnostic RADreports to enhance report accuracy further.

Limitations


The SERPENT-VLM has shown significant advancements in creating radiology reports from chest X-rays, reducing inaccuracies, and better matching the content of the images compared to earlier models. However, this research has its limitations. The testing of the model’s performance and adaptability has been limited to particular datasets (IU X-Ray and ROCO), which do not encompass the broad spectrum of radiological images or health conditions. It remains unclear how well this would work in actual medical situations. Furthermore, although the model’s ability to handle low-quality images is emphasized, the wide range of image quality in real-life scenarios could pose challenges that have yet to be evaluated.

Ethics Statement

The deployment of SERPENT-VLM in clinical settings involves significant ethical considerations. The model’s potential to generate erroneous interpretations from radiological images, despite re-



(a) Performance metrics for ROCO dataset with varying levels of Gaussian noise added to input radiological images.



(b) Performance metrics for IU-Xray dataset with varying levels of Gaussian noise added to input radiological images.

Figure 3: Comparative performance metrics for ROCO and IU-Xray datasets.

duced hallucinations, necessitates cautious application, especially since incorrect reports could lead to misdiagnoses or inappropriate treatments. The use of large datasets for training also raises privacy concerns, requiring stringent data handling and patient consent protocols.

References

2020. Language models are few-shot learners. In *Advances in Neural Information Processing Systems*, volume 33, pages 1877–1901. Curran Associates, Inc.
- Anas Awadalla, Irena Gao, Josh Gardner, Jack Hessel, Yusuf Hanafy, Wanrong Zhu, Kalyani Marathe, Yonatan Bitton, Samir Gadre, Shiori Sagawa, et al. 2023. Openflamingo: An open-source framework for training large autoregressive vision-language models. *arXiv preprint arXiv:2308.01390*.
- Zhihong Chen, Yaling Shen, Yan Song, and Xiang Wan. 2021. Cross-modal memory networks for radiology report generation. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 5904–5914, Online. Association for Computational Linguistics.
- Zhihong Chen, Yan Song, Tsung-Hui Chang, and Xiang Wan. 2020. [Generating radiology reports via memory-driven transformer](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 1439–1449, Online. Association for Computational Linguistics.
- Stephanie L Hyland, Shruthi Bannur, Kenza Bouzid, Daniel C Castro, Mercy Ranjit, Anton Schwaighofer, Fernando Pérez-García, Valentina Salvatelli, Shaury Srivastav, Anja Thieme, et al. 2023. Maira-1: A specialised large multimodal model for radiology report generation. *arXiv preprint arXiv:2311.13668*.
- Jeremy Irvin, Pranav Rajpurkar, Michael Ko, Yifan Yu, Silvana Ciurea-Ilcus, Chris Chute, Henrik Marklund, Behzad Haghgoo, Robyn Ball, Katie Shpanskaya, Jayne Seekins, David A. Mong, Safwan S. Halabi, Jesse K. Sandberg, Ricky Jones, David B. Larson, Curtis P. Langlotz, Bhavik N. Patel, Matthew P. Lungren, and Andrew Y. Ng. 2019. Chexpert: A large chest radiograph dataset with uncertainty labels and expert comparison. In *Proceedings of the Thirty-Third AAAI Conference on Artificial Intelligence and Thirty-First Innovative Applications of Artificial Intelligence Conference and Ninth AAAI Symposium on Educational Advances in Artificial Intelligence*, AAAI’19/IAAI’19/EAAI’19. AAAI Press.
- Baoyu Jing, Pengtao Xie, and Eric Xing. 2018. On the automatic generation of medical imaging reports. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*. Association for Computational Linguistics.
- Hung Le, Yue Wang, Akhilesh Deepak Gotmare, Silvio Savarese, and Steven Hoi. 2022. CodeRL: Mastering code generation through pretrained models and deep reinforcement learning. In *Advances in Neural Information Processing Systems*.
- Chunyuan Li, Cliff Wong, Sheng Zhang, Naoto Usuyama, Haotian Liu, Jianwei Yang, Tristan Naumann, Hoifung Poon, and Jianfeng Gao. 2023. [Llava-med: Training a large language-and-vision assistant for biomedicine in one day](#).
- Yehao Li, Yingwei Pan, Jingwen Chen, Ting Yao, and Tao Mei. 2021. X-modaler: A versatile and high-performance codebase for cross-modal analytics. In *Proceedings of the 29th ACM International Conference on Multimedia*, MM ’21, page 3799–3802, New York, NY, USA. Association for Computing Machinery.
- Fenglin Liu, Xian Wu, Shen Ge, Wei Fan, and Yuexian Zou. 2021. Exploring and distilling posterior and prior knowledge for radiology report generation. In *2021 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. IEEE.
- Haotian Liu, Chunyuan Li, Qingyang Wu, and Yong Jae Lee. 2024. Visual instruction tuning. *Advances in neural information processing systems*, 36.
- Farhad Nooralahzadeh, Nicolas Perez Gonzalez, Thomas Frauenfelder, Koji Fujimoto, and Michael

- Krauthammer. 2021. Progressive transformer-based generation of radiology reports. In *Findings of the Association for Computational Linguistics: EMNLP 2021*, pages 2824–2832, Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Long Ouyang, Jeffrey Wu, Xu Jiang, Diogo Almeida, Carroll Wainwright, Pamela Mishkin, Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, John Schulman, Jacob Hilton, Fraser Kelton, Luke Miller, Maddie Simens, Amanda Askell, Peter Welinder, Paul F Christiano, Jan Leike, and Ryan Lowe. 2022. Training language models to follow instructions with human feedback. In *Advances in Neural Information Processing Systems*, volume 35, pages 27730–27744. Curran Associates, Inc.
- Yingwei Pan, Ting Yao, Yehao Li, and Tao Mei. 2020. X-linear attention networks for image captioning. In *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. IEEE.
- Niket Tandon, Aman Madaan, Peter Clark, and Yiming Yang. 2022. Learning to repair: Repairing model output errors after deployment using a dynamic memory of feedback. In *Findings of the Association for Computational Linguistics: NAACL 2022*, pages 339–352, Seattle, United States. Association for Computational Linguistics.
- Mingkang Tang, Zhanyu Wang, Zhenhua LIU, Fengyun Rao, Dian Li, and Xiu Li. 2021. Clip4caption: Clip for video caption. In *Proceedings of the 29th ACM International Conference on Multimedia*, MM '21, page 4858–4862, New York, NY, USA. Association for Computing Machinery.
- Mingkang Tang, Zhanyu Wang, Zhaoyang Zeng, Xiu Li, and Luping Zhou. 2023. Stay in grid: Improving video captioning via fully grid-level representation. *IEEE Transactions on Circuits and Systems for Video Technology*, 33(7):3319–3332.
- Gemini Team, Rohan Anil, Sebastian Borgeaud, Yonghui Wu, Jean-Baptiste Alayrac, Jiahui Yu, Radu Soricut, Johan Schalkwyk, Andrew M Dai, Anja Hauth, et al. 2023. Gemini: a family of highly capable multimodal models. *arXiv preprint arXiv:2312.11805*.
- O. Vinyals, A. Toshev, S. Bengio, and D. Erhan. 2015. Show and tell: A neural image caption generator. In *2015 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 3156–3164, Los Alamitos, CA, USA. IEEE Computer Society.
- Jun Wang, Abhir Bhalerao, and Yulan He. 2022a. Cross-modal prototype driven network for radiology report generation. In *Computer Vision – ECCV 2022*, pages 563–579, Cham. Springer Nature Switzerland.
- Zhanyu Wang, Hongwei Han, Lei Wang, Xiu Li, and Luping Zhou. 2022b. Automated radiographic report generation purely on transformer: A multicriteria supervised approach. *IEEE Transactions on Medical Imaging*, 41(10):2803–2813.
- Zhanyu Wang, Luping Zhou, Lei Wang, and Xiu Li. 2021. A self-boosting framework for automated radiographic report generation. In *2021 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 2433–2442.
- Kelvin Xu, Jimmy Ba, Ryan Kiros, Kyunghyun Cho, Aaron Courville, Ruslan Salakhudinov, Rich Zemel, and Yoshua Bengio. 2015. Show, attend and tell: Neural image caption generation with visual attention. In *Proceedings of the 32nd International Conference on Machine Learning*, volume 37 of *Proceedings of Machine Learning Research*, pages 2048–2057, Lille, France. PMLR.
- An Yan, Zexue He, Xing Lu, Jiang Du, Eric Chang, Amilcare Gentili, Julian McAuley, and Chun-Nan Hsu. 2021. Weakly supervised contrastive learning for chest X-ray report generation. In *Findings of the Association for Computational Linguistics: EMNLP 2021*, pages 4009–4015, Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Michihiro Yasunaga and Percy Liang. 2020. Graph-based, self-supervised program repair from diagnostic feedback. In *Proceedings of the 37th International Conference on Machine Learning, ICML'20*. JMLR.org.
- Di You, Fenglin Liu, Shen Ge, Xiaoxia Xie, Jing Zhang, and Xian Wu. 2021. AlignTransformer: Hierarchical alignment of visual regions and disease tags for medical report generation. In *Medical Image Computing and Computer Assisted Intervention – MICCAI 2021*, Lecture notes in computer science, pages 72–82. Springer International Publishing, Cham.
- Quanzeng You, Hailin Jin, Zhaowen Wang, Chen Fang, and Jiebo Luo. 2016. Image captioning with semantic attention. In *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 4651–4659.
- Jianbo Yuan, Haofu Liao, Rui Luo, and Jiebo Luo. 2019. Automatic radiology report generation based on multi-view image fusion and medical concept enrichment. In *Medical Image Computing and Computer Assisted Intervention – MICCAI 2019*, pages 721–729, Cham. Springer International Publishing.
- Kai Zhang, Jun Yu, Eashan Adhikarla, Rong Zhou, Zhiling Yan, Yixin Liu, Zhengliang Liu, Lifang He, Brian Davison, Xiang Li, Hui Ren, Sunyang Fu, James Zou, Wei Liu, Jing Huang, Chen Chen, Yuyin Zhou, Tianming Liu, Xun Chen, Yong Chen, Quanzheng Li, Hongfang Liu, and Lichao Sun. 2024. [Biomedgpt: A unified and generalist biomedical generative pre-trained transformer for vision, language, and multimodal tasks.](#)
- Deyao Zhu, Jun Chen, Xiaoqian Shen, Xiang Li, and Mohamed Elhoseiny. 2023. [Minigpt-4: Enhancing vision-language understanding with advanced large language models.](#)