

Argus: Benchmarking and Enhancing Vision-Language Models for 3D Radiology Report Generation

Che Liu¹, Zhongwei Wan², Yuqi Wang³, Hui Shen⁴,
Haozhe Wang⁵, Kangyu Zheng⁶, Mi Zhang², Rossella Arcucci¹

¹Imperial College London ²The Ohio State University ³University of Hong Kong

⁴University of Michigan ⁵Hong Kong University of Science and Technology

⁶Rensselaer Polytechnic Institute

che.liu@imperial.ac.uk

Abstract

Automatic radiology report generation holds significant potential to streamline the labor-intensive process of report writing by radiologists, particularly for 3D radiographs such as CT scans. While CT scans are critical for clinical diagnostics, they remain less explored compared to 2D radiographs. To date, there has been no comprehensive benchmark for 3D radiograph report generation (3DRRG), nor sufficient investigation into the optimal training strategies for Vision Language Models (VLMs) in this context, particularly with respect to vision encoder choices, visual token compression, and model scaling. In this work, we make three key contributions. We curate **CT-3DRRG**, the largest **publicly** available 3D CT-report dataset, establishing a robust and diverse benchmark for evaluating VLM performance on 3DRRG. Furthermore, we propose a comprehensive training recipe for building high-performing VLMs for 3DRRG, exploring key factors such as vision encoder pretraining strategies, visual token compression, and the impact of data & model scale. Guided by these findings, we introduce **Argus**, a state-of-the-art family of VLMs that achieve superior performance across different model sizes and input 3D medical image resolutions, efficiently processing high-resolution 3D images up to $512 \times 512 \times 256$. The code will be released at <https://github.com/cheliu-computation/Argus>.

1 Introduction

Radiology reports are essential for clinical decision making, yet their manual creation is labor-intensive and time-consuming (Bastawrous and Carney, 2017; Rimmer, 2017; Rosenkrantz et al., 2016). This has driven the need for automation in radiology report generation. Existing work on this task primarily focuses on 2D images, such as chest X-rays (Tanida et al., 2023; Li et al., 2023c). However, compared to 2D images, 3D medical images

provide more comprehensive diagnostic information and are crucial for identifying life-threatening diseases such as pulmonary opacities and early-stage cancer (Bradley et al., 2019; Self et al., 2013).

While 3D radiograph report generation (3DRRG) has the potential to significantly improve clinical workflows, several challenges remain unresolved in the literature. Notably, there is no comprehensive benchmark for 3DRRG, and many existing studies lack robust evaluations across multiple datasets and metrics. Traditional NLP metrics alone are inadequate for assessing the clinical relevance of generated reports. As a result, new evaluation metrics, such as clinical efficacy measures like GREEN (Ostmeier et al., 2024) and RaTEScore (Zhao et al., 2024), are essential for inclusion in the benchmark to evaluate report quality in clinical settings.

Additionally, high-resolution (HR) 3D images offer more clinical value than low-resolution (LR) images, as certain conditions, such as pulmonary nodules, are harder to detect in LR scans (Liu et al., 2021b). However, most studies (Bai et al., 2024; Wu et al., 2023) downsample HR volumes (e.g., $512 \times 512 \times 256$) to LR ($256 \times 256 \times 32$) for report generation, which leads to significant information loss and may overlook critical details. This downsampling is primarily driven by the increased computational cost associated with processing the larger number of visual tokens generated by HR volumes. Furthermore, no clear training recipe exists for effectively building Vision-Language Models (VLMs) for 3DRRG, particularly in terms of selecting and pretraining vision encoders, visual token compression, and optimizing model and data scale.

This brings us to a crucial challenge: **the lack of large-scale public datasets and a comprehensive benchmark** for 3DRRG. Additionally, **training strategies for VLMs remain unclear**, requiring a systematic analysis of key components.

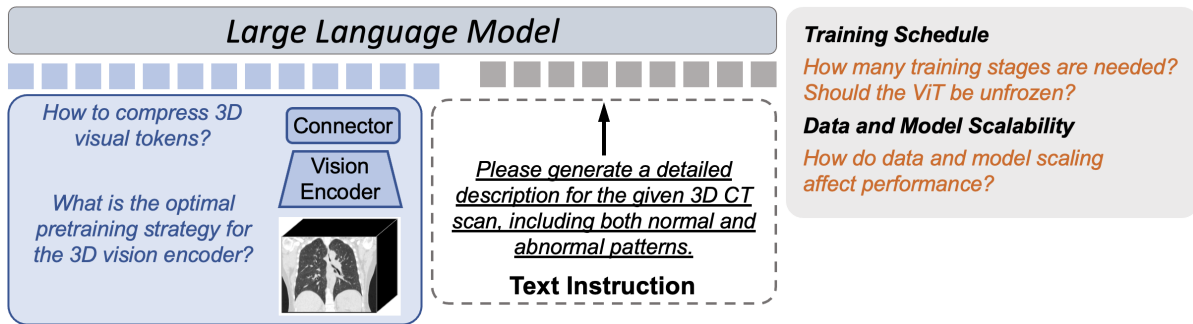


Figure 1: **Argus framework**. A schematic illustration of our comprehensive exploration of 3DRRG, covering key design choices from 3D vision encoder pretraining and visual token compression to training schedules and scalability. We systematically analyze each component in Sections 4 and 5.

To address this challenge, we make 4 key contributions:

- We curate the largest publicly available 3D CT-report dataset, **CT-3DRRG**, comprising 49,849 CT volumes paired with radiology reports from three publicly available sources: BIMCV-R, CT-RATE, and INSPECT (Chen et al., 2024a; Hamamci et al., 2024a; Huang et al., 2023).
- Based on CT-3DRRG, we establish the first comprehensive benchmark for 3DRRG, incorporating 8 NLP metrics and 3 clinical efficacy metrics. This enables broad evaluation across diverse sources, moving beyond single-resource or in-house data.
- We conduct an in-depth analysis of the key components required to build VLMs for 3DRRG, including strategies for **vision encoder pretraining** and **visual token compression**, the **VLM training schedule**, and the impact of **data and model scaling**.
- We propose a family of VLMs, **Argus**, that achieves superior performance on 3DRRG, surpassing existing methods in both NLP and clinical efficacy metrics. Argus scales from 3B to 70B parameters and efficiently handles 3D image resolutions up to $512 \times 512 \times 256$.

2 Related Works

Vision-Language Models. The development of Vision-Language Models (VLMs) combines computer vision and natural language processing to enhance visual and linguistic capabilities. This integration is crucial for tasks that require both visual perception and language comprehension. Models like CLIP (Radford et al., 2019), Flamingo (Alayrac et al., 2022), and BLIP (Li et al., 2023a)

have improved the alignment between these modalities by using extensive image-text data samples, resulting in significant performance gains. LLaVA (Liu et al., 2023) further simplifies this approach by using a basic linear projector to align visual features with language space, involving only a few learnable parts and tailored instruction data to fully leverage the model’s strong capabilities. Despite their effectiveness, these approaches focus mainly on 2D images. To explore the relatively underexplored area of 3D medical imaging, such as CT scans, we extend the LLaVA framework by incorporating a customized 3D vision encoder specifically designed for 3D image variants. Our goal is to tackle the challenge of generating accurate radiology reports for 3D medical images.

Radiology Report Generation for 2D Medical Images. Radiology report generation for 2D images typically involves template-based generation with region differentiation (Li et al., 2018; Jing et al., 2017), knowledge integration and coherence to highlight abnormalities (Liu et al., 2021a; Ma et al., 2021), and cross-modal alignment using attention to link textual and visual features (Gu et al., 2024; Tanida et al., 2023; Li et al., 2023c). Although effective for 2D settings (e.g., chest X-rays), extending these methods to 3D volumetric data (e.g., CT scans) is challenging. The 3D context involves distinct patterns in each slice (e.g., lungs and heart in upper slices versus abdominal organs in lower slices), whereas 2D images typically present uniform spatial structures with intensity variations (Adegun et al., 2021; Puttagunta and Ravi, 2021; Singh et al., 2020).

Radiology Report Generation for 3D Medical Images. Despite its clinical significance, only a few studies (Hamamci et al., 2024b; Wu et al., 2023; Bai et al., 2024) have attempted 3D radiograph report generation (3DRRG) with publicly available code, while others (Lai et al., 2024; Chen




Figure 2: Distribution of the CT-3DRRG dataset. **Left:** Word cloud visualization highlighting the most frequent terms in the radiology reports. **Right:** Histogram showing the distribution of report lengths (measured in number of tokens).

et al., 2024b) do not release their code, hindering direct comparisons. Current methods have notable drawbacks: CT2Rep (Hamamci et al., 2024b) processes high-resolution (HR) volumes slice by slice, losing 3D context; RadFM (Wu et al., 2023) and M3D (Bai et al., 2024) down-sample HR inputs to low-resolution (LR) views and use Preceiver (Alayrac et al., 2022) to compress large numbers of visual tokens to just 64 or apply spatial pooling, potentially corrupting critical visual information. Preserving HR is crucial because downsampling coronal/sagittal axes can obscure subtle abnormalities (Bradley et al., 2019; Self et al., 2013), and clinical practice often favors 256-slice CT scans for enhanced diagnostic detail (Klass et al., 2010; Chua et al., 2013). Moreover, these approaches rely on small-scale or permission-restricted datasets, limiting their broader evaluation and generalizability.

To advance 3DRRG, we curate the largest publicly accessible 3D CT-report dataset for evaluating VLM performance. We also conduct a thorough investigation of crucial components in VLM training, enabling our approach to process 3D scans at up to $512 \times 512 \times 256$ resolution without downsampling, thus preserving essential diagnostic details.

3 Benchmark Construction

3.1 Curating the CT-3DRRG Dataset

To establish a robust and diverse benchmark for 3D Radiology Report Generation (3DRRG), we construct the CT-3DRRG dataset by integrating three publicly available resources: BIMCV-R (Chen et al., 2024a), CT-RATE (Hamamci et al., 2024a), and INSPECT (Huang et al., 2023). To ensure data quality, we apply a systematic curation process to the original datasets, as described in Section A.

Dataset	Train/Val/Test Split	Avg. Tokens per Report	Max Tokens per Report
BIMCV-R	3,726/532/1,064	97.19	406
CT-RATE	21,715/2,412/1,564	201.40	824
INSPECT	14,280/2,040/4,080	48.71	267
Total	39,721/4,984/5,144	130.03	824

Table 1: Processed data details of CT-3DRRG.

Notably, we retain the entire CT-RATE official test set without filtering. After curation, the number of samples in each dataset is summarized in Table 1. The BIMCV-R and INSPECT datasets are split into training, validation, and test sets following a 70%/10%/20% split. For the CT-RATE dataset, we utilize its official test set and allocate 10% of the training set as a validation set. Thus, the CT-3DRRG dataset consists of three subsets corresponding to these sources.

Following curation, we propose CT-3DRRG, a unified dataset constructed by merging multiple publicly available datasets. The statistical details of CT-3DRRG are presented in Table 1 and Figure 2. During training, we combine the training sets from all sources into a single dataset. For evaluation, we evaluate VLM performance on the test set of each subset separately, allowing for a thorough analysis of robustness and generalizability across different data distributions.

3.2 Data Preprocessing

After building the CT-3DRRG dataset, we preprocess the data for training and evaluating the VLM. The preprocessing steps are designed to standardize the data and ensure it is in the appropriate format for model training.

HU Value Clipping and Intensity Normalization: The Hounsfield Unit (HU) values of the CT scans are clipped to the range of -1000 to +1000 to

standardize the intensity levels. Then, the intensity values are normalized to the $[0, 1]$ range to ensure uniform scaling across all images.

Spatial Normalization: Following the RadGenome-Chest CT (Zhang et al., 2024), the spacing of the CT scans is normalized to $[1, 1, 4]$ mm on the sagittal, coronal, and transverse axes to ensure consistency in voxel spacing across samples.

Resizing: All CT scans are resized to two different resolution settings. For high-resolution settings, the scans are resized to $512 \times 512 \times 256$, where 512 represents the dimensions of the sagittal and coronal axes, and 256 is the number of slices in the transverse axis. For normal-resolution settings, the scans are resized to $256 \times 256 \times 64$, where 256 refers to the size in the sagittal and coronal axes, and 64 is the number of slices in the transverse axis.

To preserve the integrity of spatial relationships, which is crucial for accurate interpretation of 3D medical images, augmentation techniques such as flipping or rotating were not employed. This decision was made to avoid potential ambiguities in spatial orientation, particularly given radiological reports that specify locations such as the “left lung” and “right lung.”

For the instructions used during training the VLM, we utilize the instruction: “*Please generate a detailed description for the given 3D CT scan, including both normal and abnormal patterns.*”. Inspired by (Bai et al., 2024) where diverse instructions are employed to prevent overfitting on a specific instruction style. To further enhance diversity, we use GPT-4o to rephrase the instruction into 30 distinct variants, all maintaining the same semantic meaning.

3.3 Selecting Appropriate Metrics

For the 3DRRG task, the most straightforward evaluation approach involves using NLP-based metrics such as BLEU, METEOR, CIDEr, and ROUGE (Bai et al., 2024). However, these metrics primarily assess lexical and syntactic similarity while overlooking the clinical correctness and relevance of the generated reports. Some studies (Hamamci et al., 2024a,c) attempt to address this limitation by employing text classifiers, such as those in CT2Rep (Hamamci et al., 2024c), to evaluate report accuracy. However, these classifiers are restricted to limited predefined categories, making them unsuitable for more diverse

datasets like BIMCV-R (Chen et al., 2024a) and INSPECT (Huang et al., 2023). To provide a clinically meaningful assessment, we adopt RadGraphXL (Delbrouck et al., 2024), GREEN (Ostmeier et al., 2024), and RaTEScore (Zhao et al., 2024) as evaluation metrics. These metrics offer a more reliable evaluation by focusing on structured clinical information, factual consistency, and radiological correctness, leveraging specifically trained models for clinical assessment.

4 Exploring VLM Design for 3DRRG

4.1 Vision Encoder Pre-training

In this section, we investigate the pre-training strategy of the vision encoder for the VLM. We use the ViT-Base (Dosovitskiy et al., 2020) architecture, incorporating 3D positional embeddings and patching to adapt it for 3D CT volumes. To assess the impact of resolution, we explore two settings: for normal-resolution CT scans ($256 \times 256 \times 64$), we use a patch size of $16 \times 16 \times 8$, generating 2048 visual tokens; for high-resolution scans ($512 \times 512 \times 256$), a patch size of $32 \times 32 \times 16$ produces 4096 tokens. These configurations allow us to analyze how resolution influences the encoder’s ability to learn meaningful visual representations. To comprehensively evaluate the impact of different pre-training strategies on the vision encoder, we explore a range of visual self-supervised learning (vSSL) methods and language supervision techniques.

- **3D-MAE (Zhou et al., 2023):** Masks visual tokens to reconstruct voxel intensities, learning visual representations.
- **Vox2Vec (Goncharov et al., 2023):** Aims to learn consistent voxel-level representations across views, enhancing voxel-wise representation consistency.
- **FLIP (Li et al., 2023b):** Pretrains the visual encoder with patch masking, aligning visual representations with textual descriptions.
- **Vox2Vec + FLIP:** Trains in two stages: (1) learns voxel-wise representations with Vox2Vec, (2) applies contrastive learning to align 3D CT with reports.
- **3D-MAE + FLIP:** Trains in two stages: (1) reconstructs masked patches with 3D-MAE, (2) applies contrastive learning to align 3D CT with reports.




Figure 3: Comparison of different model combinations involving various vision encoder pretraining strategies, visual token compression, and connector designs under (a) normal resolution and (b) high resolution settings. The results represent the average clinical metric across the three subsets mentioned in Section 3.

For methods involving patch masking, we set the masking ratio to 50%. Additionally, we analyze different masking ratios in Figure 5 to assess their impact on model performance.

4.2 Visual Token Compression and Connector

In this section, we explore various strategies for visual token compression and the corresponding connectors, as the large number of visual tokens generated from 3D CT scans increases the cost of visual token processing in the VLM. In M3D (Bai et al., 2024), a 3D average pooling is applied to the ViT output visual tokens, and MLP is used as a connector to convert the visual token dimensions to the LLM embedding dimensions. Meanwhile, RadFM (Wu et al., 2023) uses a Preceiver to compress these tokens into 64 randomly initialized learnable queries. However, these methods all result in some loss of visual information. To address this, we introduce pixel shuffle (Shi et al., 2016; Team, 2024) as a non-information-loss compression method to investigate the optimal combination of visual token compression and connectors. For both pooling and pixel shuffle, we set the down-sampling factor to 0.5 on each dimension. This means that for normal resolution, 2048 tokens are downsampled to 256 tokens, and for high resolution, 4096 tokens are downsampled to 512 tokens.

As shown in Figure 3, we combine these token compression strategies on the features extracted from the vision encoder, pretrained with various methods mentioned in Section 4.1. The results show that the 2-layer MLP with pixel shuffle, using a vision encoder pretrained with 3D-MAE + FLIP, outperforms the other methods across all metrics. This provides the following key findings:

1. **3D-MAE combined with language supervision achieves superior performance.** This is




Figure 4: Comparison of different training schedules, including multi-stage strategies and variations of ViT, under (a) normal resolution and (b) high resolution settings. The results reflect the average clinical metric computed across the three subsets outlined in Section 3.

because 3D-MAE does not rely on data augmentation techniques, unlike Vox2Vec (Goncharov et al., 2023), which uses rotation or flipping that can distort the spatial relationships in 3D CT scans. Additionally, the two-stage visual pretraining helps progressively enhance the model’s ability to learn more accurate and meaningful representations.

2. **Pixel shuffle is more effective than other compression methods, as it does not result in information loss.** Pixel shuffle reduces the number of tokens while increasing the dimensionality of each token, retaining all visual information. In contrast, pooling performs brute-force averaging of all pixels, which can lead to a loss of spatial detail. The Preceiver method, using random query initialization to compress all visual tokens into only 64 queries, results in substantial information loss, as discussed in (Yao et al., 2024). Additionally, the 2-layer MLP with various compression methods consistently outperforms the 1-layer MLP.

Based on these insights, we recommend that the

optimal VLM should train the vision encoder using visual SSL combined with language supervision, utilize pixel shuffle for visual token compression, and adopt a 2-layer MLP as the connector strategy.

5 How to Properly Train a VLM for 3DRRG?

5.1 Training Schedules

For our VLM training, we select Llama3.1-instruct as the LLM backbone for the 8B and 70B models, and Llama3.2-instruct for the 3B models. For the vision encoder and token compression strategies, we follow the findings presented in Sections 4.1 and 4.2.

In this section, we explore the impact of different training schedules. In the **1-stage** approach, we jointly train the connector and the LLM for one epoch using supervised fine-tuning (SFT). The **2-stage** approach consists of two steps: first, we freeze the ViT and LLM while training only the connector for one epoch; second, we train the connector, MLP, and LLM together using SFT for one epoch. We use the DeepSpeed ZeRO-3 offload setting for all model scales to ensure consistency in our experiments. The hyperparameters are detailed in Section B

As depicted in Figure 4, we find that the 2-stage approach consistently outperforms the 1-stage method, aligning with the findings in (Tong et al., 2024). Additionally, we investigate the effect of freezing versus unfreezing the ViT. Our experiments indicate that unfreezing the ViT consistently enhances performance, particularly under the high-resolution setting, where it achieves the highest performance. This suggests that allowing the ViT to update its weights refines feature representations, leading to improved 3DRRG performance. Given these findings, we recommend using the 2-stage training approach with an unfrozen ViT to further boost performance.

5.2 Investigating the Effects of Data and Model Scales

In this section, we explore how scaling both data and model size affects VLM training performance. As shown in Figure 6, we begin with a dataset of 10K samples and progressively increase the size to 20K, 30K, and finally the full dataset. The results demonstrate a consistent improvement in performance as the dataset size increases, highlighting the model’s scalability. We evaluate this effect

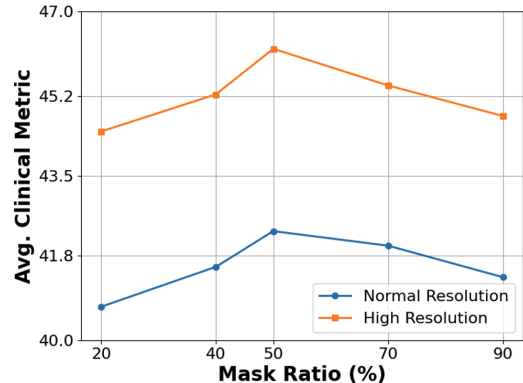


Figure 5: Ablation study on the masking ratio for vision encoder pretraining. We evaluate the impact of the mask ratio with 3D-MAE combined with the FLIP strategy using a 8B LLM.

on both normal-resolution and high-resolution settings.

Additionally, we examine the impact of scaling the LLM size from 3B to 70B parameters, as shown in Figure 6. The results indicate that larger models yield better performance, further validating the scalability of our training strategy. These findings suggest that our approach is effective and can seamlessly scale to accommodate both larger datasets and more powerful LLMs.

6 Argus, the State-of-the-Art VLM for 3DRRG

We leverage insights from our studies to develop Argus, a state-of-the-art (SOTA) VLM for the 3DRRG task. To ensure a fair comparison, we evaluate Argus alongside three existing baselines: RadFM (Wu et al., 2023), M3D (Bai et al., 2024), and CT2Rep (Hamamci et al., 2024c), all of which have publicly available code, allowing for direct evaluation.

To maintain consistency, we replace the LLM backbones of these baselines with Llama3.1-Instruct (8B and 70B) and Llama3.2-Instruct-3B, aligning them with our LLM selection. We reimplement their training on our CT-3DRRG dataset while using the official vision encoder weights from their repositories. Additionally, we adhere to their original implementations for token compression strategies, connectors, and training schedules. For the high-resolution setting, we interpolate the positional embeddings of their vision encoders to support a resolution of $512 \times 512 \times 256$, enabling them to process high-resolution 3D CT scans. We compare Argus across both normal ($256 \times 256 \times 64$)

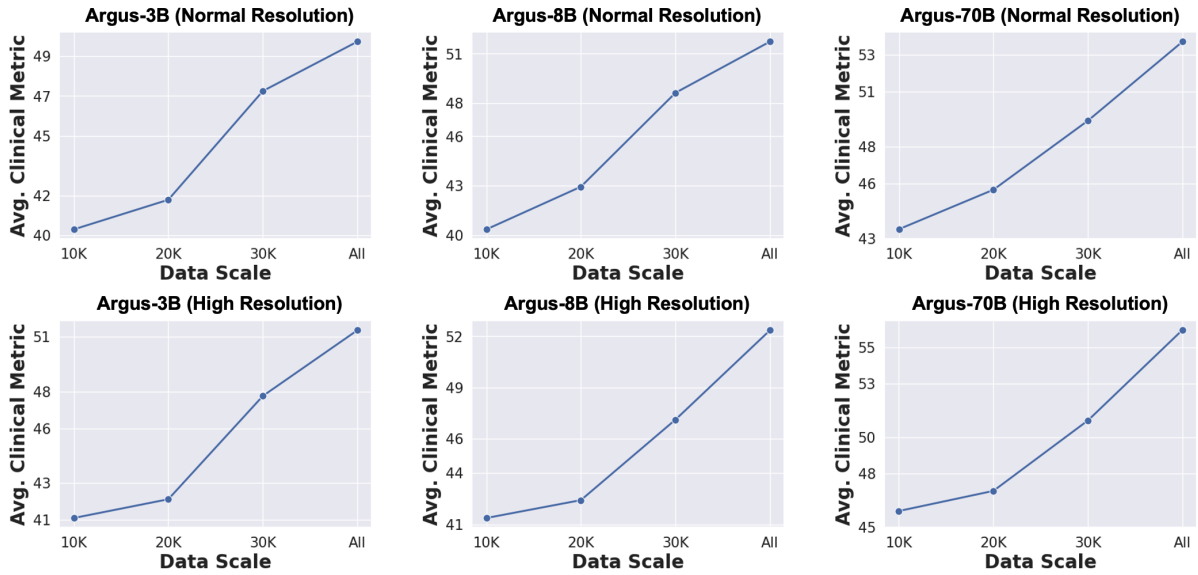


Figure 6: Scaling study of Argus with different LLM sizes and data scales. We evaluate Argus at 3B, 8B, and 70B parameters under both normal and high-resolution settings while increasing the data scale from 10K to the full dataset. The results show a consistent improvement in performance as both model size and data scale increase, demonstrating the robustness and scalability of our approach.

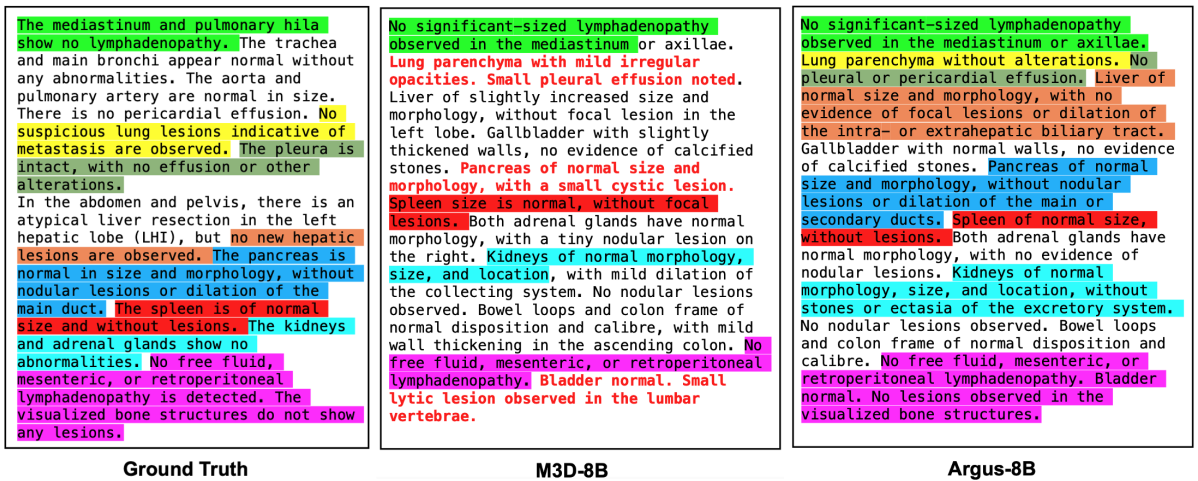


Figure 7: Qualitative comparison between the ground truth and reports generated by existing methods (M3D-8B) and our method (Argus-8B). Highlighted background text indicates correctly generated content, while bold red text denotes incorrect information. M3D-8B exhibits inconsistencies in lesion description and misidentifies key anatomical structures, whereas Argus-8B generates more accurate and clinically relevant descriptions, demonstrating its superiority in medical report generation.

and high-resolution ($512 \times 512 \times 256$) settings.

All models are trained on the same training set from CT-3DRRG and evaluated on three test sets. We report performance using average NLP metrics (BLEU1-4, ROUGE-1,2,L, METEOR, CIDEr) and clinical evaluation metrics, including GREEN, RadGraphXL, and RaTEScore (Ostmeier et al., 2024; Delbrouck et al., 2024; Zhao et al., 2024), as shown in Table 2. Argus consistently outperforms the baselines across the 3B-70B model range and both resolution settings. Notably, Argus-3B surpasses or matches the performance of other 8B models, while

Argus-8B achieves comparable results to the 70B models of other methods. At both resolution settings, Argus outperforms baselines with the same LLM size.

We further visualize the generated reports from M3D-8B and Argus-8B under the normal resolution setting, as shown in Figure 7. The results demonstrate that Argus-8B accurately captures key anatomical structures and pathological findings, while M3D-8B exhibits inconsistencies in lesion descriptions and misidentifies certain structures. This highlights the superiority of our method in

Model	CT-RATE				BIMCV-R				INSPECT			
	Avg. NLP	GREEN	RadT _{EScore}	RadGraphXL	Avg. NLP	GREEN	RadT _{EScore}	RadGraphXL	Avg. NLP	GREEN	RadT _{EScore}	RadGraphXL
<i>Normal Resolution Setting (256×256×64)</i>												
CT2Rep-3B	38.77	48.96	58.10	25.14	30.95	43.65	37.18	19.74	27.55	36.62	38.10	24.93
RadFM-3B	39.95	49.46	59.02	25.73	31.48	44.10	37.56	20.15	28.09	37.23	39.01	25.36
M3D-3B	40.23	50.31	60.16	26.08	32.10	44.94	38.55	20.11	28.82	37.96	39.62	25.41
Argus-3B	41.53	54.07	64.82	28.12	33.52	47.12	40.83	21.28	30.14	40.03	41.86	27.52
CT2Rep-8B	40.52	51.72	62.14	26.45	32.15	45.15	39.19	20.27	29.55	38.50	41.18	26.34
RadFM-8B	40.61	52.21	62.78	27.14	32.60	45.57	40.55	20.72	29.89	39.09	41.65	26.68
M3D-8B	41.33	52.80	63.20	26.95	33.32	46.02	40.11	21.15	30.33	39.54	41.71	27.15
Argus-8B	43.31	56.60	67.67	29.82	35.02	48.79	42.23	22.16	31.77	41.48	43.40	28.71
CT2Rep-70B	44.02	55.10	64.42	28.15	33.50	47.01	43.12	21.54	31.40	40.04	42.44	27.90
RadFM-70B	44.40	55.67	65.10	28.35	34.18	48.32	42.65	21.74	31.33	40.90	43.12	28.14
M3D-70B	44.95	56.12	65.32	30.10	35.05	47.75	43.09	22.11	32.10	41.67	43.95	28.95
Argus-70B	45.83	58.24	70.02	31.07	36.42	50.33	44.57	23.28	33.04	42.72	45.51	30.03
<i>High Resolution Setting (512×512×256)</i>												
CT2Rep-3B	41.05	51.02	60.53	27.83	32.68	46.12	38.45	21.49	29.62	38.92	40.64	26.53
RadFM-3B	41.78	51.52	61.24	28.24	33.24	46.49	39.06	21.86	30.06	39.42	41.31	26.88
M3D-3B	42.32	52.18	62.02	28.66	34.02	47.07	39.77	22.03	30.64	39.83	42.16	27.26
Argus-3B	43.03	55.56	66.53	29.52	35.11	48.82	42.33	22.72	31.74	41.53	43.59	28.93
CT2Rep-8B	43.70	53.10	66.12	29.57	35.54	49.31	42.02	22.63	32.48	41.12	43.66	29.21
RadFM-8B	44.03	53.62	66.64	30.07	36.00	49.83	42.44	23.14	32.84	41.59	44.21	29.66
M3D-8B	44.75	54.23	67.20	30.62	36.48	50.10	43.11	23.39	33.38	42.05	44.58	30.05
Argus-8B	45.81	58.42	70.68	32.04	37.07	51.27	44.58	24.19	33.82	43.47	45.92	30.72
CT2Rep-70B	46.35	56.25	68.20	30.27	36.15	49.60	45.01	23.60	33.35	42.62	45.00	30.66
RadFM-70B	46.72	57.36	68.85	30.56	36.62	50.22	45.39	23.77	33.70	43.09	44.55	30.04
M3D-70B	47.26	57.12	69.34	31.55	37.15	50.64	46.81	24.22	34.15	43.53	45.34	30.62
Argus-70B	47.98	60.22	72.53	33.07	38.62	52.43	47.05	25.53	35.22	45.02	47.72	32.31

Table 2: Performance comparison of Argus across different scales on three sub-test sets of CT-3DRRG. We evaluate models ranging from 3B to 70B parameters and compare them with CT2Rep (Hamamci et al., 2024b), RadFM (Wu et al., 2023), and M3D (Bai et al., 2024) under both normal and high-resolution settings. Argus-3B outperforms other 8B models and even achieves comparable performance with larger-scale models. Argus-8B matches the performance of other 70B variants, while at the 70B scale, Argus surpasses all competing methods. These results demonstrate the superiority of our approach, highlighting its scalability and effectiveness across resolutions.

generating clinically precise and coherent radiology reports. We also implement human expert evaluation on all method results as shown in Section C, where the expert feedback corroborates our findings, further emphasizing the advantages of Argus-8B in terms of clinical accuracy and report quality. Argus achieves SOTA performance and efficiency across model scales, demonstrating the effectiveness of our architectural and training strategies without the need for larger model sizes.

7 Conclusion

In this work, we present **CT-3DRRG**, the first and largest dataset comprising 3D volume-report pairs for 3DRRG, built entirely from **publicly** available resources. We critically examine key components

of VLMs, including vision encoder pre-training, visual token compression, and associated design and training strategies, while also analyzing the impact of scaling both data and model sizes. Leveraging these findings, we introduce **Argus**, a family of VLMs ranging from 3B to 70B parameters, which achieves superior performance on 3DRRG task. Thanks to our efficient design, even smaller models such as Argus-3B outperform existing methods utilizing larger LLMs (e.g., 7B), demonstrating the effectiveness and efficiency of our approach. Our study provides valuable insights into how these factors influence VLM performance in 3DRRG tasks and establishes a strong foundation for future research, fostering the development of more efficient and scalable VLMs for 3DRRG.

Limitation

In this work, we design a comprehensive benchmark for 3DRRG and investigate various components of VLMs, but due to the costly nature of human evaluation, generated reports cannot be evaluated manually for every report. Additionally, human evaluation is inherently difficult to quantify, which poses challenges in establishing clear metrics for report quality. Moreover, due to data scarcity and computational resource limitations, scaling the dataset to the millions, as seen in natural image domains, remains a significant challenge. In the future, we aim to develop more efficient evaluation metrics that can automate the quality assessment process and explore strategies for scaling both the dataset and model more effectively.

References

- Adekanmi Adeyinka Adegun, Serestina Viriri, and Rose-line Oluwaseun Ogundokun. 2021. Deep learning approach for medical image analysis. *Computational Intelligence and Neuroscience*, 2021:1–9.
- Jean-Baptiste Alayrac, Jeff Donahue, Pauline Luc, Antoine Miech, Iain Barr, Yana Hasson, Karel Lenc, Arthur Mensch, Katherine Millican, Malcolm Reynolds, et al. 2022. Flamingo: a visual language model for few-shot learning. *Advances in Neural Information Processing Systems*, 35:23716–23736.
- Fan Bai, Yuxin Du, Tiejun Huang, Max Q-H Meng, and Bo Zhao. 2024. M3d: Advancing 3d medical image analysis with multi-modal large language models. *arXiv preprint arXiv:2404.00578*.
- Sarah Bastawrous and Benjamin Carney. 2017. Improving patient safety: avoiding unread imaging exams in the national va enterprise electronic health record. *Journal of digital imaging*, 30:309–313.
- Stephen H Bradley, Sarah Abraham, Matthew EJ Callister, Adam Grice, William T Hamilton, Rocio Rodriguez Lopez, Bethany Shinkins, and Richard D Neal. 2019. Sensitivity of chest x-ray for detecting lung cancer in people presenting with symptoms: a systematic review. *British Journal of General Practice*, 69(689):e827–e835.
- Yinda Chen, Che Liu, Xiaoyu Liu, Rossella Arcucci, and Zhiwei Xiong. 2024a. Bimcv-r: A landmark dataset for 3d ct text-image retrieval. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 124–134. Springer.
- Zhixuan Chen, Yequan Bie, Haibo Jin, and Hao Chen. 2024b. Large language model with region-guided referring and grounding for ct report generation. *arXiv preprint arXiv:2411.15539*.
- Su-Kiat Chua, Hwei-Fong Hung, Jun-Jack Cheng, Min-Tsung Tseng, Wai-Yip Law, Chu-Jen Kuo, Chiung-Zuan Chiu, Che-Ming Chang, Shih-Huang Lee, Huey-Ming Lo, et al. 2013. Diagnostic performance of 64-versus 256-slice computed tomography coronary angiography compared with conventional coronary angiography in patients with suspected coronary artery disease. *Acta Cardiologica Sinica*, 29(2):151.
- Jean-Benoit Delbrouck, Pierre Chambon, Zhihong Chen, Maya Varma, Andrew Johnston, Louis Blanke-meier, Dave Van Veen, Tan Bui, Steven Truong, and Curtis Langlotz. 2024. Radgraph-xl: A large-scale expert-annotated dataset for entity and relation extraction from radiology reports. In *Findings of the Association for Computational Linguistics ACL 2024*, pages 12902–12915.
- Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, et al. 2020. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929*.
- Mikhail Goncharov, Vera Soboleva, Anvar Kurmukov, Maxim Pisov, and Mikhail Belyaev. 2023. vox2vec: A framework for self-supervised contrastive learning of voxel-level representations in medical images. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 605–614. Springer.
- Tiancheng Gu, Dongnan Liu, Zhiyuan Li, and Weidong Cai. 2024. Complex organ mask guided radiology report generation. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 7995–8004.
- Ibrahim Ethem Hamamci, Sezgin Er, Furkan Almas, Ayse Gulnihhan Simsek, Sevval Nil Esirgun, Irem Dogan, Muhammed Furkan Dasdelen, Bastian Wittmann, Enis Simsar, Mehmet Simsar, et al. 2024a. A foundation model utilizing chest ct volumes and radiology reports for supervised-level zero-shot detection of abnormalities. *arXiv preprint arXiv:2403.17834*.
- Ibrahim Ethem Hamamci, Sezgin Er, and Bjoern Menze. 2024b. Ct2rep: Automated radiology report generation for 3d medical imaging. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 476–486. Springer.
- Ibrahim Ethem Hamamci, Sezgin Er, and Bjoern Menze. 2024c. Ct2rep: Automated radiology report generation for 3d medical imaging. *arXiv preprint arXiv:2403.06801*.
- Shih-Cheng Huang, Zepeng Huo, Ethan Steinberg, Chia-Chun Chiang, Matthew P Lungren, Curtis P Langlotz, Serena Yeung, Nigam H Shah, and Jason A Fries. 2023. Inspect: a multimodal dataset for pulmonary embolism diagnosis and prognosis. *arXiv preprint arXiv:2311.10798*.
- Baoyu Jing, Pengtao Xie, and Eric Xing. 2017. On the automatic generation of medical imaging reports. *arXiv preprint arXiv:1711.08195*.
- Oliver Klass, Matthew Walker, Anne Siebach, Tina Stuber, Sebastian Feuerlein, Markus Juchems, and Martin HK Hoffmann. 2010. Prospectively gated axial ct coronary angiography: comparison of image quality and effective radiation dose between 64- and 256-slice ct. *European radiology*, 20:1124–1131.
- Haoran Lai, Zihang Jiang, Qingsong Yao, Rongsheng Wang, Zhiyang He, Xiaodong Tao, Wei Wei, Weifu Lv, and S Kevin Zhou. 2024. E3d-gpt: Enhanced 3d visual foundation for medical vision-language model. *arXiv preprint arXiv:2410.14200*.
- Junnan Li, Dongxu Li, Silvio Savarese, and Steven Hoi. 2023a. Blip-2: Bootstrapping language-image pre-training with frozen image encoders and large language models. *arXiv preprint arXiv:2301.12597*.

- Yanghao Li, Haoqi Fan, Ronghang Hu, Christoph Feichtenhofer, and Kaiming He. 2023b. Scaling language-image pre-training via masking. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 23390–23400.
- Yaowei Li, Bang Yang, Xuxin Cheng, Zhihong Zhu, Hongxiang Li, and Yuexian Zou. 2023c. Unify, align and refine: Multi-level semantic alignment for radiology report generation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 2863–2874.
- Yuan Li, Xiaodan Liang, Zhiting Hu, and Eric P Xing. 2018. Hybrid retrieval-generation reinforced agent for medical image report generation. *Advances in neural information processing systems*, 31.
- Fenglin Liu, Xian Wu, Shen Ge, Wei Fan, and Yuexian Zou. 2021a. Exploring and distilling posterior and prior knowledge for radiology report generation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 13753–13762.
- Haotian Liu, Chunyuan Li, Qingyang Wu, and Yong Jae Lee. 2023. Visual instruction tuning. *arXiv preprint arXiv:2304.08485*.
- Shurong Liu, Hongbo Liu, Peipei Li, and Lijie Jiang. 2021b. Application of high-resolution ct images information in complicated infection of lung tumors. *Journal of Infection and Public Health*, 14(3):418–422.
- Xuwei Ma, Fenglin Liu, Changchang Yin, Xian Wu, Shen Ge, Yuexian Zou, Ping Zhang, and Xu Sun. 2021. Contrastive attention for automatic chest x-ray report generation. *arXiv preprint arXiv:2106.06965*.
- Sophie Ostmeier, Justin Xu, Zhihong Chen, Maya Varma, Louis Blankemeier, Christian Bluethgen, Arne Edward Michalson, Michael Moseley, Curtis Langlotz, Akshay S Chaudhari, et al. 2024. Green: Generative radiology report evaluation and error notation. *arXiv preprint arXiv:2405.03595*.
- Muralikrishna Puttagunta and S Ravi. 2021. Medical image analysis based on deep learning approach. *Multi-media tools and applications*, 80(16):24365–24398.
- Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, Ilya Sutskever, et al. 2019. Language models are unsupervised multitask learners. *OpenAI blog*, 1(8):9.
- Abi Rimmer. 2017. Radiologist shortage leaves patient care at risk, warns royal college. *BMJ: British Medical Journal (Online)*, 359.
- Andrew B Rosenkrantz, Danny R Hughes, and Richard Duszak Jr. 2016. The us radiologist workforce: an analysis of temporal and geographic variation by using large national datasets. *Radiology*, 279(1):175–184.
- Wesley H Self, D Mark Courtney, Candace D McNaughton, Richard G Wunderink, and Jeffrey A Kline. 2013. High discordance of chest x-ray and computed tomography for detection of pulmonary opacities in ed patients: implications for diagnosing pneumonia. *The American journal of emergency medicine*, 31(2):401–405.
- Wenzhe Shi, Jose Caballero, Ferenc Huszár, Johannes Totz, Andrew P Aitken, Rob Bishop, Daniel Rueckert, and Zehan Wang. 2016. Real-time single image and video super-resolution using an efficient sub-pixel convolutional neural network. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1874–1883.
- Satya P Singh, Lipo Wang, Sukrit Gupta, Haveesh Goli, Parasuraman Padmanabhan, and Balázs Gulyás. 2020. 3d deep learning on medical images: a review. *Sensors*, 20(18):5097.
- Tim Tanida, Philip Müller, Georgios Kaissis, and Daniel Rueckert. 2023. Interactive and explainable region-guided radiology report generation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7433–7442.
- OpenGVLab Team. 2024. Internvl2: Better than the best—expanding performance boundaries of open-source multimodal models with the progressive scaling strategy.
- Shengbang Tong, Ellis Brown, Penghao Wu, Sanghyun Woo, Manoj Middepogu, Sai Charitha Akula, Jihan Yang, Shusheng Yang, Adithya Iyer, Xichen Pan, et al. 2024. Cambrian-1: A fully open, vision-centric exploration of multimodal llms. *arXiv preprint arXiv:2406.16860*.
- Chaoyi Wu, Xiaoman Zhang, Ya Zhang, Yanfeng Wang, and Weidi Xie. 2023. Towards generalist foundation model for radiology. *arXiv preprint arXiv:2308.02463*.
- Linli Yao, Lei Li, Shuhuai Ren, Lean Wang, Yuanxin Liu, Xu Sun, and Lu Hou. 2024. Deco: Decoupling token compression from semantic abstraction in multimodal large language models. *arXiv preprint arXiv:2405.20985*.
- Xiaoman Zhang, Chaoyi Wu, Ziheng Zhao, Jiayu Lei, Ya Zhang, Yanfeng Wang, and Weidi Xie. 2024. Radgenome-chest ct: A grounded vision-language dataset for chest ct analysis. *arXiv preprint arXiv:2404.16754*.
- Weike Zhao, Chaoyi Wu, Xiaoman Zhang, Ya Zhang, Yanfeng Wang, and Weidi Xie. 2024. Ratescore: A metric for radiology report generation. *arXiv preprint arXiv:2406.16845*.
- Lei Zhou, Huidong Liu, Joseph Bae, Junjun He, Dimitris Samaras, and Prateek Prasanna. 2023. Self pre-training with masked autoencoders for medical image classification and segmentation. In *2023 IEEE 20th International Symposium on Biomedical Imaging (ISBI)*, pages 1–6. IEEE.

A Details of CT-3DRRG Dataset Curation

The CT-3DRRG dataset is curated from three public datasets: BIMCV-R (Chen et al., 2024a), CT-RATE (Hamamci et al., 2024a), and INSPECT (Huang et al., 2023). Each of these datasets contains paired 3D CT scans and corresponding radiology reports. The curation process involves several key steps to ensure the dataset’s quality and relevance for training and evaluation.

- **CT-RATE Dataset:** In the CT-RATE dataset (Hamamci et al., 2024a), each sample contains both the ‘FINDINGS’ and ‘IMPRESSION’ sections of the radiology report. These two sections are concatenated into a single report.
- **BIMCV-R and INSPECT Datasets:** These datasets only contain one part of the report. For these, we directly use the available report without modification.
- **Duplicate CT Scans in CT-RATE:** The CT-RATE dataset includes duplicate CT scans for each sample, generated through different reconstruction methods. To maintain consistency with the RadGenome-Chest CT (Zhang et al., 2024), we retain only one CT scan per sample and pair it with the corresponding report.

Following dataset preparation, we implement several filtering and cleaning steps to ensure the quality of the reports:

- **Removal of Numerical Values Not Directly Obtainable from CT Scans:** Using GPT-4o, we filter sentences that contain numerical values that cannot be directly inferred from the CT scans. For example, sentences such as “SAT O2 without oxygen of 93” and “Fever up to 38” are removed, as these values are typically obtained from the patient’s electronic health records (EHR) rather than from the CT scan.
- **Removal of Numerical Measurements Dependent on External Tools:** We also remove sentences containing numerical values that require external tools for precise measurement, such as “Increase in trunk caliber of the 39 mm pulmonary artery”. Since the goal of the

3DRRG dataset is to generate patterns rather than measure exact sizes, such sentences are discarded.

- **Removal of Sentences Comparing to Previous Studies:** Sentences referencing previous studies, such as “It is compared to the previous study of March 2019, not mediastinal,” are also removed. These comparisons are impractical as not all samples are paired with prior CT scans. Furthermore, our work aims to build a benchmark for 3DRRG using single CT scans, rather than longitudinal data.
- **Minimum Report Length:** To ensure the relevance of the reports, we remove samples with reports containing fewer than 10 tokens.

After these curation steps, the dataset consists of the following number of samples for each source dataset: 5,322 samples for BIMCV-R, 25,691 samples for CT-RATE, and 20,400 samples for INSPECT. These samples are then split into training, validation, and test sets, as shown in Table 1.

As a result of these curation efforts, we have created the largest publicly available dataset for 3D Radiology Report Generation (3DRRG), which we have named **CT-3DRRG**. This dataset is entirely based on public resources, ensuring both its accessibility and relevance to the broader research community. The CT-3DRRG dataset is now a valuable resource for advancing research in the field of automated radiology report generation from 3D CT scans.

B Hyper-parameters of Training

In this study, we implement supervised fine-tuning (SFT) for the 3DRRG task over two epochs, conducting all experiments on eight A100-80G GPUs. For computational efficiency, we employ mixed precision with `bf16`, as detailed in Table 3. The training process uses a total batch size of 8 with a gradient accumulation factor of 2, resulting in an effective batch size of 16. The maximum sequence length is restricted to 1024 tokens.

Regarding learning rates, we use 1×10^{-4} during the first stage, where only the connector is updated, and 1×10^{-6} during the second stage, where the ViT, connector, and LLM are updated. For the 3B and 8B models, we update all LLM parameters, while for the 70B model, we apply LoRA with an alpha value of 64, a rank of 128, and a dropout

Table 3: Hyperparameters for our method.

Hyperparameter	Value
Mixed precision	bf16
Total epochs	1
Total effective batch size	16
Gradient accumulation	1
Maximum sequence length	1024
Learning rate (Stage 1)	1×10^{-4}
Learning rate (Stage 2)	1×10^{-6}
Optimizer	AdamW
Schedule	Linear
Warm-up ratio	0.05
Weight decay	0.0

rate of 0.1 to reduce memory consumption while maintaining fine-tuning efficiency.

Our models are optimized using the AdamW algorithm with a linear learning rate decay, a warm-up ratio of 0.05, and a weight decay of 0.0. We utilize the same configuration for all existing methods unless their official codebase specifies different hyperparameters.

C Human Evaluation of Generated Reports

Model	CT-RATE	BIMCV	INSPECT
CT2Rep-7B	2.9 ± 1.3	2.7 ± 1.4	2.6 ± 1.1
RadFM-7B	3.2 ± 0.9	3.0 ± 1.3	2.8 ± 1.3
M3D-7B	3.3 ± 1.1	3.1 ± 0.9	2.8 ± 1.2
Argus-7B	4.1 ± 0.7	3.7 ± 0.8	3.4 ± 1.0

Table 4: Comparison of CT-RATE, BIMCV, and INSPECT scores for different models.

We conducted a human evaluation of the model-generated reports from CT2Rep-7B, RadFM-7B, M3D-7B, and Argus-7B under normal resolution setting. We generated 100 reports for each test subset and asked three human experts to evaluate the correctness of the generated reports on a scale of 1 to 5. The results are shown in Table 4. As shown, Argus-7B achieved the highest score, demonstrating its superior performance in generating accurate and reliable reports.