

Gla-AI4BioMed at RRG24: Visual Instruction-tuned Adaptation for Radiology Report Generation

Xi Zhang, Zaiqiao Meng*, Jake Lever and Edmond S. L. Ho

School of Computing Science, University of Glasgow

X.Zhang.6@research.gla.ac.uk, Zaiqiao.Meng@glasgow.ac.uk

Jake.Lever@glasgow.ac.uk, Shu-Lim.Ho@glasgow.ac.uk

Abstract

We introduce a radiology-focused visual language model designed to generate radiology reports from chest X-rays. Building on previous findings that large language models (LLMs) can acquire multimodal capabilities when aligned with pretrained vision encoders, we demonstrate similar potential with chest X-ray images. This integration enhances the ability of model to understand and describe chest X-ray images. Our model combines an image encoder with a fine-tuned LLM based on the Vicuna-7B architecture, enabling it to generate different sections of a radiology report with notable accuracy. The training process involves a two-stage approach: (i) initial alignment of chest X-ray features with the LLM (ii) followed by fine-tuning for radiology report generation¹.

1 Introduction

Radiology reports constitute the primary medium through which radiologists convey the findings and conclusions derived from radiography, such as chest X-rays. These reports play a pivotal role in the diagnostic and therapeutic processes across a wide range of diseases, emphasizing their significance in contemporary medical practice (Engle et al., 2021). Structured to enhance clarity and efficacy in medical communication, radiology reports primarily feature FINDINGS and IMPRESSIONS sections (Kahn et al., 2009). The FINDINGS section details the critical observations of the radiologist on the image, while the IMPRESSIONS section summarizes the conclusions and recommendations of the radiologist. These sections collectively ensure that radiology reports are indispensable in diagnostic and therapeutic decision-making, combining image analysis and clinical insight. Table 1 shows an example generated by GPT-4 (OpenAI et al., 2024), which delineates these sections.

* Corresponding author.

¹<https://github.com/Glasgow-AI4BioMed/RRG-BioNLP-ACL2024>.

FINDINGS

There has been an increase in size of the left pleural effusion compared to the prior exam. The right lung remains clear with no evidence of consolidation or pneumothorax. The heart size is mildly enlarged but stable. The mediastinum appears unremarkable. Mild degenerative changes are noted in the thoracic spine and ribs. The upper abdomen is without remarkable findings.

IMPRESSIONS

Increase in left pleural effusion compared to prior. Stable mild cardiomegaly. No evidence of right lung pathology.

Table 1: FINDINGS and IMPRESSIONS in a synthetic radiology report generated by GPT-4.

Radiology report generation (RRG) is crucial for advancing future medical artificial intelligence systems (Monshi et al., 2020). This task involves transforming images into text, necessitating alignment between imaging and textual data. Significant advancements in natural language processing have driven progress in this area, with large generative visual language models like LLaVA (Liu et al., 2023), InstructBLIP (Dai et al., 2023), and Flamingo (Alayrac et al., 2022) leading the way.

The prevailing visual language models, such as those mentioned above, aim to address the challenge of multimodal alignment by leveraging large-scale pretraining. Typically, this involves adapting a vision encoder for integration with a pretrained LLM. To meet specific task requirements, various degrees of finetuning are applied. For example, LLaVA (Liu et al., 2023) represents a novel end-to-end trained large multimodal model for general-purpose visual and language understanding, achieving impressive chat capabilities. However, in the context of our work, the focus is on fine-tuning

image-text pairs, specifically for tasks related to medical images, to enhance the capability of the visual language model in radiology report generation.

In the specialized area of radiographic report generation, it is paramount for models to discern nuanced details within multiple medical images. These details include subtle variations in opacity against a backdrop of overlapping structures (Panayides et al., 2020). Therefore, the radiographic report generation task extends beyond the mere extraction of details from a single image. Models must interpret the clinical implications of these nuances to generate precise and medically rigorous text in reports. This is a particularly crucial capability for radiographic report generation models, since it enhances the clinical utility and effectiveness of radiology reports and ensures their accuracy and relevance in clinical settings.

General-domain models have proven inadequate for generating findings in radiology reports (Hyland et al., 2024). In this work, we propose a radiology-specific visual language model designed for solving the radiology report generation task by fine-tuning across various sections of medical reports. Our model utilizes a two-stage fine-tuning process that significantly enhances its performance. In particular, we initially align the large language model with image embedding through a pretraining phase. In the second stage, we further fine-tune the LLM using Low-Rank adaptation (LoRA) techniques (Hu et al., 2021). Both stages are trained on the dataset in this workshop (Xu et al., 2024).

Additionally, we use a straightforward strategy of merging and stitching multiple images to form a single cohesive input, enabling the model to effectively process and integrate information from multiple X-ray images. Using the dataset provided by this workshop, which includes a collection of chest X-rays and their corresponding sections, we fine-tune our model to enhance the accuracy and specificity of the generated radiology reports.

This paper investigates the fine-tuning of visual instruction for a visual language model in the specific domain of radiology report generation. We describe the training of two distinct models developed for the Shared Task on Large-Scale Radiology Report Generation (RRG24) at the BioNLP 2024 Workshop (Xu et al., 2024). In the public test set, we achieved an F1-RadGraph score (Delbrouck et al., 2022a) of 24.13 and 22.79 in the Findings and Impressions sections, respectively.

In the hidden test set, we achieved F1-RadGraph scores (Delbrouck et al., 2022a) of 24.13 and 22.10 in the Findings and Impressions sections, respectively, which places us 4th on the leaderboard at the time of submission. The contributions of this research are as follows:

- We enhance domain adaptation for radiology by implementing visual instruction tuning, which further fine-tunes the visual language model specifically for image-to-text tasks. This approach optimizes performance in interpreting and translating visual data into descriptive, clinically relevant text.
- We adopt a method of stitching multiple images together, allowing a single image encoder to process multiple image inputs simultaneously. This strategy obviates the need for separate encoding of each image, enabling the model to adapt to varying numbers of image inputs using limited resources.

2 Related Work

Nowadays, exemplified by open-source projects such as LLaVA (Liu et al., 2023), the effectiveness of self-supervised vision-language models (VLMs) using parallel data has been demonstrated in different research domains. These VLMs, when instruction-tuned with multimodal inputs, align well with human intentions and perform robustly in various downstream tasks, including converting images to text (Park and Kim, 2023).

However, the unique characteristics of biomedical image-text pairs significantly differ from those in general domains. Biomedical images often contain subtle and complex features that require precise interpretation, while the corresponding text must convey highly specific medical information (Huff et al., 2021). In biomedical settings, VLMs designed for general domains often fail to meet these specialized needs, as they lack the ability to accurately interpret medical data and generate relevant clinical descriptions (Chang et al., 2023). This discrepancy underscores the urgent need for domain-specific fine-tuning. By tailoring VLMs to the distinct demands of the biomedical field, such fine-tuning can enhance their ability to capture and convey the intricate details necessary for accurate medical interpretations and reports.

Recent advancements have been made in adapting general-purpose foundation models for medical applications, particularly in radiology. The

Med-Flamingo (Moor et al., 2023), an extension of the OpenFlamingo framework (Awadalla et al., 2023), leverages images and captions from medical textbooks to enhance few-shot visual question-answering capabilities. Similarly, the Med-PaLM M, developed by Tu et al. (2023), fine-tuned the PaLM-E model (Driess et al., 2023) using comprehensive biomedical datasets. LLaVA-Med, proposed by Li et al. (2023), modifies the LLaVA (Liu et al., 2023) framework with image-text pairings and multimodal instructions from PubMed data. Additionally, the ELIXR model, developed by Xu et al. (2023), integrates the SupCon CXR encoder (Sellergren et al., 2022) with the PaLM 2-S model (Anil et al., 2023) to support classification, semantic search, question answering, and quality assurance. Finally, the Radiology-GPT, created by Liu et al. (2024), utilizes radiology reports from MIMIC-CXR (Johnson et al., 2019) to facilitate the generation of findings-to-impression text, based on the Alpaca instruction-tuning framework (Taori et al., 2023).

Historically, research in radiology report generation has varied, with some studies focusing exclusively on either the Findings or the Impressions sections (Jin et al., 2024; Yan et al., 2023), while others have addressed both. Notably, Endo et al. (2021) and Bannur et al. (2023) specialized in generating only the Impressions section. In contrast, studies by Miura et al. (2021), Delbrouck et al. (2022a), Tanida et al. (2023), Nicolson et al. (2023), and Tu et al. (2023) concentrated on the Findings section. Comprehensive analyses by Yu et al. (2023) and Jeong et al. (2023) covered all settings, demonstrating that the choice of sections significantly influences reported performance metrics, complicating comparative evaluations across different study designs.

However, these existing models have limitations. Most notably, they are typically designed to process single images and often fall short in generating reports that match the depth and detail of those written by human radiologists. Additionally, they do not fully replicate the workflow of medical professionals, who often reference multiple images to enhance report accuracy. Our work addresses these gaps by developing a model capable of handling multiple images simultaneously and generating comprehensive radiology reports. This approach aims to more closely mimic the process used by medical professionals, thereby improving the accuracy and quality of the generated reports.

3 Methodology

In our study, we follow the observations from LLaVA-Med (Li et al., 2023), suggesting superior performance when initiating with a language-only pretrained LLM rather than a multimodal-trained base. Our model architecture incorporates an image encoder and a learnable adapter placed atop the image outputs, mirroring the LLaVA-1.5 model design (Liu et al., 2023). We adopt an auto-regressive language modelling approach using cross-entropy loss (Graves, 2014) and align hyperparameters with those from LLaVA-1.5, including a joint tuning phase for the LLM and adapter (Liu et al., 2023). In alignment with LLaVA-1.5 protocols, we initially pretrain the adapter alone for one epoch, followed by a full training cycle lasting three epochs, employing Low-Rank Adaptation of Large Language Models techniques (LoRA) (Hu et al., 2021) for efficient parameter tuning.

3.1 Task Description

A key application of natural language generation in medicine is developing support systems that produce written reports from X-ray images, detailing clinical findings. Such systems are highly valuable, potentially reducing the routine workload of radiologists and improving the efficacy of clinical interactions. The objective of this shared task is to generate radiology reports from one or more chest X-rays taken during a single study, specifically targeting two sections: ‘Findings’ and ‘Impressions’ (as shown in Table 1).

Consequently, our team is dedicated to the task of exclusively producing either the ‘Findings’ or ‘Impressions’ sections of the report. We have developed separate models for each section because their focuses are different. The ‘Findings’ section provides a factual description based on the images, while the ‘Impressions’ section offers the radiologist’s conclusions and recommendations. By separating the models, we can tailor each to better address its specific requirements.

For the radiology report section generation, handling multiple images is crucial as it allows the model to provide a detailed and accurate description of the observed facts, similar to how radiologists analyze multiple images to form a comprehensive understanding. This approach enhances the model’s ability of to mimic the actual workflow of medical professionals, who often reference multiple images.

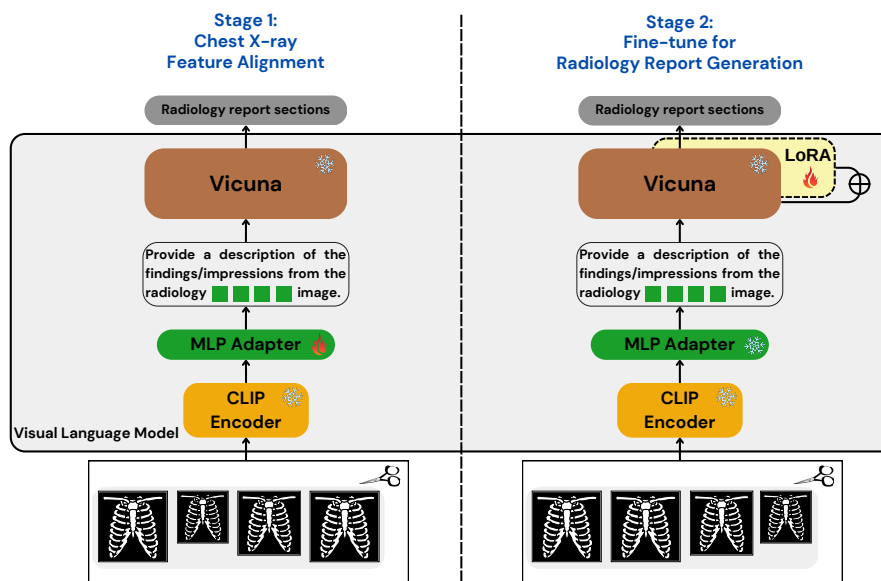


Figure 1: Our two-stage training framework. In the first stage, visual features are aligned with LLM. In the second stage, the model focuses on the training of radiology report generation tasks.

3.2 The Proposed Models

To solve the RRG24 shared task, we fine-tune a large visual language model for radiology report generation based on the provided dataset. Specifically, we propose two separate models, called Med-CXRGen-F and Med-CXRGen-I, fine-tuned on Findings and Impressions sections respectively.

We use CLIP (Radford et al., 2021) as an image encoder and Vicuna-1.5 (Chiang et al., 2023) as a large language model. Our adaptation module consists of a multi-layer perceptron (MLP) featuring GELU activations (Hendrycks and Gimpel, 2023) and a uniform hidden size of 1024 across all layers.

The interaction with the model involves alternating system messages linked with the corresponding image. The training objective of the model is to generate accurate responses. Initially, we convert the image into a series of image patch tokens via the image encoder, selecting embeddings from the penultimate layer. These image features are then processed by the MLP adapter, aligning them to the input specifications of the LLM.

The instructional prompt of the report generation task we employed is: "Provide a description of the findings/impressions from the radiology <image>\n image." In this prompt, "<image>\n" represents the image holder token, as shown in Figure 1, which indicates to the LLM that it should

base its generation on the input image.

3.3 Training

The same network architecture is utilized for different radiology report sections, where an MLP adapter connects the vision encoder and the language model. For model training, we use a two-stage procedure: (as shown in Figure 1)

- **Stage 1: Chest X-ray Feature Alignment**

In the first epoch training phase on the provided dataset, each sample, accompanied by instructions and image input, prompts the model to predict the original caption. During this stage, we keep the visual encoder and LLM weights unchanged, focusing solely on updating the MLP adapter. This approach aligns the features from chest X-ray images with their textual embeddings in the LLM. Training is limited to a single epoch, which facilitates the expansion of the vocabulary of aligned image-text tokens specific to the radiology domain.

- **Stage 2: Fine-tune for Radiology Report Generation**

In the second phase, the visual encoder weights and adapter are kept frozen while continuing to update the pre-trained LLM weights using LoRA (Hu et al., 2021) technology. Further fine-tuning is conducted on the provided dataset through visual instrumental tuning with three epochs.

4 Evaluation

4.1 Dataset

We fine-tune and evaluate our models using the RRG24 dataset hosted on the BioNLP ACL’24 (Xu et al., 2024), which includes data from MIMIC-CXR (Johnson et al., 2019), CheXpert (Chambron et al., 2024), PadChest (Bustos et al., 2020), BIMCV-COVID19 (de la Iglesia Vayá et al., 2020), and OpenI, with their statistics shown in Table 2.

Dataset	FINDINGS	IMPRESSIONS
training	344,394	366,413
validation	8,839	9,331
test-public	2,692	2,967
test-hidden	1,063	1,428

Table 2: Distribution of shared task on Large-Scale Radiology Report Generation.

We conducted training in two stages (refer to section 3.3). To ensure consistency between training and inference processes, we analysed the word count distribution, as shown in Table 3. Consequently, we have set a maximum length of 1024 for both the text input and inference output to minimise computational expense. On the other hand, as illustrated in Table 4, some datasets contain multiple images, therefore, we select up to the first four images for the image input. We merge multiple images horizontally to form a single-image input, which is proven to be robust in our experiments.

Dataset	FINDINGS	IMPRESSIONS
training	259 (± 180)	216 (± 153)
validation	257 (± 176)	217 (± 155)
test-public	380 (± 161)	257 (± 224)

Table 3: Average word count and standard deviation on Large-Scale Radiology Report Generation.

Dataset	FINDINGS	IMPRESSIONS
training	1.57 (± 0.63)	1.45 (± 0.62)
validation	1.58 (± 0.62)	1.45 (± 0.62)
test-public	1.70 (± 0.71)	1.67 (± 0.71)

Table 4: Average number of images and standard deviation on Large-Scale Radiology Report Generation.

4.2 Metrics

We assess the generated reports through a dual approach involving both general lexical metrics and specialized radiology metrics. Focusing on the accuracy of described medical findings, radiology-specific metrics provide a deeper insight into the clinical relevance of the reports, beyond surface-level phrasing variations. According to the RRG24 guidelines, we consider five evaluation metrics for this work, including BLEU4 (Papineni et al., 2002), ROUGEL (Lin, 2004), BERT score (Zhang et al., 2020), F1-cheXbert (Smit et al., 2020), and F1-RadGraph (Delbrouck et al., 2022a).

4.3 Training details

We evaluated our two proposed models, i.e. Med-CXRGen-F and Med-CXRGen-I, on the workshop evaluation datasets, based on a computational infrastructure utilizing an A6000 GPU (48GB memory each) with the Deepspeed zero-3 configuration (Rajbhandari et al., 2020) with BF16 enabled. We employ a cosine learning rate scheduler that begins with a warm-up phase of 0.03 and sets the learning rate at $1 \cdot 10^{-5}$. The global batch size for our experiments is set at 16. Observations of the smallest loss on the evaluation dataset throughout the training process guide us to select this as the final checkpoint for all runs. For inference on the test dataset, we decode in 32-bit precision up to 150 tokens, consistent with the baseline model on the leaderboard (Delbrouck et al., 2022b). Each model required approximately 215 hours of training.

5 Results

We report the performance of our two proposed models over five evaluation metrics in Table 5. As shown in Table 5, in the public test set, we achieved an F1-RadGraph score (Delbrouck et al., 2022a) of 24.13 and 22.79 in the Findings and Impressions sections, respectively. In the hidden test set, we achieved F1-RadGraph scores (Delbrouck et al., 2022a) of 24.13 and 22.10 in the Findings and

Model	Dataset	Section	BLEU4	ROUGEL	Bertscore	F1-cheXbert	F1-RadGraph
Med-CXRGen-F	validation	Findings	7.02	23.33	48.93	40.42	21.94
	test-public	Findings	8.07	24.90	53.45	45.91	24.13
	test-hidden	Findings	7.65	24.35	52.69	46.21	24.13
Med-CXRGen-I	validation	Impressions	10.18	28.10	51.78	50.51	26.65
	test-public	Impressions	7.10	25.11	47.39	47.43	22.79
	test-hidden	Impressions	9.60	25.27	48.60	46.74	22.10

Table 5: Evaluation results on different datasets.

Impressions sections, respectively, which places us 4th on the leaderboard² at the time of submission. Additionally, our model achieved notable Bertscore results in the test-public set, with 53.45 for Findings and 47.39 for Impressions. These results demonstrate the effectiveness of our approach in generating high-quality medical reports across different datasets.

6 Discussion

Performance disparities observed between the Findings and Impressions sections of the test results can be attributed to several factors. Firstly, the Impression and Findings sections address distinct medical purposes, resulting in performance disparities. The Findings section offers an objective description of symptoms, while the Impressions are oriented towards diagnostic conclusions. The variability in word count between these sections also affects the complexity of model inference, as reflected in the lexical evaluation scores.

Additionally, significant discrepancies in the medical evaluation metrics highlight a varied distribution of diseases within the test set. This heterogeneity could impact the generalisability and accuracy of the model. Furthermore, our analysis indicates that the performance may be compromised in multi-image inference scenarios where it does not account for superfluous images. Such factors are essential to consider when assessing the diagnostic accuracy and reliability of the model in clinical settings. Enhancing the ability of model to differentiate between relevant and superfluous images could significantly improve diagnostic accuracy.

Furthermore, exploring domain-specific adaptations and fine-tuning strategies tailored to the unique characteristics of medical data could further enhance model performance. Incorporating temporal dynamics into the model to capture changes

over time and developing more sophisticated frameworks for generating multi-modal radiology reports are other promising avenues for future research. These advancements are expected to enhance both the practicality and accuracy of our model within clinical scenarios.

7 Conclusion

In this work, we have developed a vision-language model capable of processing multiple images. Through visual instruction tuning, we achieved alignment between two modalities and further fine-tuning for specific downstream tasks. Notably, our system attained a commendable fourth-place standing across four diverse test datasets at the RRG24 at BioNLP 2024 workshop (Xu et al., 2024), substantiating the practicality of vision-language models within specialized medical tasks.

Moving forward, we intend to conduct in-depth research into more sophisticated methods for generating multi-modal radiology reports. This will involve incorporating temporal dynamics and developing frameworks specifically focused on text generation. Such advancements are expected to enhance both the practicality and accuracy of our model within the clinical scenario.

²<https://vilmedic.app/misc/bionlp24/leaderboard>

8 Limitations

The following section outlines the limitations identified in our study:

1. **Prevalence of certain conditions:** Some diseases are more easily detected, which may lead to artificially high medical assessment scores.
2. **Imaging modalities and anatomical structures:** There is a notable imbalance in the imaging modalities and anatomical structures covered in the training dataset. Variations such as the number of images per patient and the considerable disparity in the length of medical reports exacerbate this imbalance.
3. **Radiologist and radiology department preferences:** Preferences and writing styles vary among radiologists and radiology departments. This diversity adds complexity to medical reports by introducing inconsistencies and uncertainties that are, to a certain extent, human-induced. For example, the dataset provided in this workshop demonstrates that even the same radiology section descriptions have varying styles. These elements significantly complicate the task of report generation.

These limitations highlight areas for improvement and the need for methodological refinements to enhance model effectiveness and reliability in clinical environments. These challenges were not addressed within the scope of this workshop.

Ethics Statement

The concepts and methodologies presented in this paper are based on experimental research findings. They are not currently available for other use. The data used for the training processes were exclusively sourced from the provided dataset, which complies with ethical standards regarding Patient Health Information. Adherence to these standards guarantees the responsible handling of sensitive data throughout our research.

Acknowledgments

We extend our gratitude to the BioNLP 2024 RRG24 Shared Task organizers for providing the baseline pipeline ViLMedic (Delbrouck et al., 2022b) and curating these challenging and exciting tasks (Xu et al., 2024).

References

- Jean-Baptiste Alayrac, Jeff Donahue, Pauline Luc, Antoine Miech, Iain Barr, Yana Hasson, Karel Lenc, Arthur Mensch, Katie Millican, Malcolm Reynolds, Roman Ring, Eliza Rutherford, Serkan Cabi, Tengda Han, Zhitao Gong, Sina Samangooei, Marianne Monteiro, Jacob Menick, Sebastian Borgeaud, Andrew Brock, Aida Nematzadeh, Sahand Sharifzadeh, Mikolaj Binkowski, Ricardo Barreira, Oriol Vinyals, Andrew Zisserman, and Karen Simonyan. 2022. [Flamingo: a visual language model for few-shot learning](#). *Preprint*, arXiv:2204.14198.
- Rohan Anil, Andrew M. Dai, Orhan Firat, Melvin Johnson, Dmitry Lepikhin, Alexandre Passos, Siamak Shakeri, Emanuel Taropa, Paige Bailey, Zhifeng Chen, Eric Chu, Jonathan H. Clark, Laurent El Shafey, Yanping Huang, Kathy Meier-Hellstern, Gaurav Mishra, Erica Moreira, Mark Omernick, Kevin Robinson, Sebastian Ruder, Yi Tay, Kefan Xiao, Yuanzhong Xu, Yujing Zhang, Gustavo Hernandez Abrego, Junwhan Ahn, Jacob Austin, Paul Barham, Jan Botha, James Bradbury, Siddhartha Brahma, Kevin Brooks, Michele Catasta, Yong Cheng, Colin Cherry, Christopher A. Choquette-Choo, Aakanksha Chowdhery, Clément Crepy, Shachi Dave, Mostafa Dehghani, Sunipa Dev, Jacob Devlin, Mark Díaz, Nan Du, Ethan Dyer, Vlad Feinberg, Fangxiaoyu Feng, Vlad Fienber, Markus Freitag, Xavier Garcia, Sebastian Gehrmann, Lucas Gonzalez, Guy Gur-Ari, Steven Hand, Hadi Hashemi, Le Hou, Joshua Howland, Andrea Hu, Jeffrey Hui, Jeremy Hurwitz, Michael Isard, Abe Ittycheriah, Matthew Jagielski, Wenhao Jia, Kathleen Kenealy, Maxim Krikun, Sneha Kudugunta, Chang Lan, Katherine Lee, Benjamin Lee, Eric Li, Music Li, Wei Li, YaGuang Li, Jian Li, Hyeontaek Lim, Hanzhao Lin, Zhongtao Liu, Frederick Liu, Marcello Maggioni, Aroma Mahendru, Joshua Maynez, Vedant Misra, Maysam Moussalem, Zachary Nado, John Nham, Eric Ni, Andrew Nystrom, Alicia Parrish, Marie Pellat, Martin Polacek, Alex Polozov, Reiner Pope, Siyuan Qiao, Emily Reif, Bryan Richter, Parker Riley, Alex Castro Ros, Aurko Roy, Brennan Saeta, Rajkumar Samuel, Renee Shelby, Ambrose Slone, Daniel Smilkov, David R. So, Daniel Sohn, Simon Tokumine, Dasha Valter, Vijay Vasudevan, Kiran Vodrahalli, Xuezhi Wang, Pidong Wang, Zirui Wang, Tao Wang, John Wieting, Yuhuai Wu, Kelvin Xu, Yunhan Xu, Linting Xue, Pengcheng Yin, Jiahui Yu, Qiao Zhang, Steven Zheng, Ce Zheng, Weikang Zhou, Denny Zhou, Slav Petrov, and Yonghui Wu. 2023. [Palm 2 technical report](#). *Preprint*, arXiv:2305.10403.
- Anas Awadalla, Irena Gao, Josh Gardner, Jack Hessel, Yusuf Hanafy, Wanrong Zhu, Kalyani Marathe, Yonatan Bitton, Samir Gadre, Shiori Sagawa, Jenia Jitsev, Simon Kornblith, Pang Wei Koh, Gabriel Ilharco, Mitchell Wortsman, and Ludwig Schmidt. 2023. [Openflamingo: An open-source framework for training large autoregressive vision-language models](#). *Preprint*, arXiv:2308.01390.

Shruthi Bannur, Stephanie Hyland, Qianchu Liu, Fer-

- nando Pérez-García, Maximilian Ilse, Daniel C. Castro, Benedikt Boecking, Harshita Sharma, Kenza Bouzid, Anja Thieme, Anton Schwaighofer, Maria Wetscherek, Matthew P. Lungren, Aditya Nori, Javier Alvarez-Valle, and Ozan Oktay. 2023. [Learning to exploit temporal structure for biomedical vision-language processing](#). *Preprint*, arXiv:2301.04558.
- Aurelia Bustos, Antonio Pertusa, Jose-Maria Salinas, and Maria de la Iglesia-Vayá. 2020. [Padchest: A large chest x-ray image dataset with multi-label annotated reports](#). *Medical Image Analysis*, 66:101797.
- Pierre Chambon, Jean-Benoit Delbrouck, Thomas Sounack, Shih-Cheng Huang, Zhihong Chen, Maya Varma, Steven QH Truong, Chu The Chuong, and Curtis P. Langlotz. 2024. [Chexpert plus: Augmenting a large chest x-ray dataset with text radiology reports, patient demographics and additional image formats](#). *Preprint*, arXiv:2405.19538.
- Yupeng Chang, Xu Wang, Jindong Wang, Yuan Wu, Linyi Yang, Kaijie Zhu, Hao Chen, Xiaoyuan Yi, Cunxiang Wang, Yidong Wang, Wei Ye, Yue Zhang, Yi Chang, Philip S. Yu, Qiang Yang, and Xing Xie. 2023. [A survey on evaluation of large language models](#). *Preprint*, arXiv:2307.03109.
- Wei-Lin Chiang, Zhuohan Li, Zi Lin, Ying Sheng, Zhanghao Wu, Hao Zhang, Lianmin Zheng, Siyuan Zhuang, Yonghao Zhuang, Joseph E. Gonzalez, Ion Stoica, and Eric P. Xing. 2023. [Vicuna: An open-source chatbot impressing gpt-4 with 90%* chatgpt quality](#).
- Wenliang Dai, Junnan Li, Dongxu Li, Anthony Tiong, Junqi Zhao, Weisheng Wang, Boyang Li, Pascale Fung, and Steven Hoi. 2023. [InstructBLIP: Towards general-purpose vision-language models with instruction tuning](#). In *Thirty-seventh Conference on Neural Information Processing Systems*.
- Maria de la Iglesia Vayá, Jose Manuel Saborit, Joaquim Angel Montell, Antonio Pertusa, Aurelia Bustos, Miguel Cazorla, Joaquin Galant, Xavier Barber, Domingo Orozco-Beltrán, Francisco García-García, Marisa Caparrós, Germán González, and Jose María Salinas. 2020. [Bimcv covid-19+: a large annotated dataset of rx and ct images from covid-19 patients](#). *Preprint*, arXiv:2006.01174.
- Jean-Benoit Delbrouck, Pierre Chambon, Christian Bluethgen, Emily Tsai, Omar Almusa, and Curtis Langlotz. 2022a. [Improving the factual correctness of radiology report generation with semantic rewards](#). In *Findings of the Association for Computational Linguistics: EMNLP 2022*, pages 4348–4360, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.
- Jean-benoit Delbrouck, Khaled Saab, Maya Varma, Sabri Eyuboglu, Pierre Chambon, Jared Dunnmon, Juan Zambrano, Akshay Chaudhari, and Curtis Langlotz. 2022b. [ViLMedic: a framework for research at the intersection of vision and language in medical AI](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics: System Demonstrations*, pages 23–34, Dublin, Ireland. Association for Computational Linguistics.
- Danny Driess, Fei Xia, Mehdi S. M. Sajjadi, Corey Lynch, Aakanksha Chowdhery, Brian Ichter, Ayzaan Wahid, Jonathan Tompson, Quan Vuong, Tianhe Yu, Wenlong Huang, Yevgen Chebotar, Pierre Sermanet, Daniel Duckworth, Sergey Levine, Vincent Vanhoucke, Karol Hausman, Marc Toussaint, Klaus Greff, Andy Zeng, Igor Mordatch, and Pete Florence. 2023. [Palm-e: An embodied multimodal language model](#). *Preprint*, arXiv:2303.03378.
- Daisuke Endo, Ryota Kobayashi, Ramon Bartolo, Bruno B Averbeck, Yasuko Sugase-Miyamoto, Kazuko Hayashi, Kenji Kawano, Barry J Richmond, and Shigeru Shinomoto. 2021. [A convolutional neural network for estimating synaptic connectivity from spike trains](#). *Scientific Reports*, 11(1):12087.
- Ryann L Engle, David C Mohr, Sally K Holmes, Marjorie Nealon Seibert, Melissa Afable, Jenniffer Leyson, and Mark Meterko. 2021. [Evidence-based practice and patient-centered care: doing both well](#). *Health care management review*, 46(3):174–184.
- Alex Graves. 2014. [Generating sequences with recurrent neural networks](#). *Preprint*, arXiv:1308.0850.
- Dan Hendrycks and Kevin Gimpel. 2023. [Gaussian error linear units \(gelus\)](#). *Preprint*, arXiv:1606.08415.
- Edward J. Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, and Weizhu Chen. 2021. [Lora: Low-rank adaptation of large language models](#). *Preprint*, arXiv:2106.09685.
- Daniel T Huff, Amy J Weisman, and Robert Jeraj. 2021. [Interpretation and visualization techniques for deep learning models in medical imaging](#). *Physics in Medicine & Biology*, 66(4):04TR01.
- Stephanie L. Hyland, Shruthi Bannur, Kenza Bouzid, Daniel C. Castro, Mercy Ranjit, Anton Schwaighofer, Fernando Pérez-García, Valentina Salvatelli, Shaury Srivastav, Anja Thieme, Noel Codella, Matthew P. Lungren, Maria Teodora Wetscherek, Ozan Oktay, and Javier Alvarez-Valle. 2024. [Maira-1: A specialised large multimodal model for radiology report generation](#). *Preprint*, arXiv:2311.13668.
- Jaehwan Jeong, Katherine Tian, Andrew Li, Sina Hartung, Fardad Behzadi, Juan Calle, David Osayande, Michael Pohlen, Subathra Adithan, and Pranav Rajpurkar. 2023. [Multimodal image-text matching improves retrieval-based chest x-ray report generation](#). *Preprint*, arXiv:2303.17579.
- Haibo Jin, Haoxuan Che, Yi Lin, and Hao Chen. 2024. [Promptmrg: Diagnosis-driven prompts for medical report generation](#). *Preprint*, arXiv:2308.12604.

- Alistair EW Johnson, Tom J Pollard, Seth J Berkowitz, Nathaniel R Greenbaum, Matthew P Lungren, Chihying Deng, Roger G Mark, and Steven Horng. 2019. [Mimic-cxr, a de-identified publicly available database of chest radiographs with free-text reports](#). *Scientific data*, 6(1):317.
- Charles E. Kahn, Curtis P. Langlotz, Elizabeth S. Burnside, John A. Carrino, David S. Channin, David M. Hovsepian, and Daniel L. Rubin. 2009. [Toward best practices in radiology reporting](#). *RADIOLOGY*, 252(3):852–856.
- Chunyu 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](#). *Preprint*, arXiv:2306.00890.
- Chin-Yew Lin. 2004. [ROUGE: A package for automatic evaluation of summaries](#). In *Text Summarization Branches Out*, pages 74–81, Barcelona, Spain. Association for Computational Linguistics.
- Haotian Liu, Chunyu Li, Qingyang Wu, and Yong Jae Lee. 2023. [Visual instruction tuning](#). *Preprint*, arXiv:2304.08485.
- Zhengliang Liu, Aoxiao Zhong, Yiwei Li, Longtao Yang, Chao Ju, Zihao Wu, Chong Ma, Peng Shu, Cheng Chen, Sekeun Kim, Haixing Dai, Lin Zhao, Lichao Sun, Dajiang Zhu, Jun Liu, Wei Liu, Dinggang Shen, Xiang Li, Quanzheng Li, and Tianming Liu. 2024. [Radiology-gpt: A large language model for radiology](#). *Preprint*, arXiv:2306.08666.
- Yasuhide Miura, Yuhao Zhang, Emily Tsai, Curtis Langlotz, and Dan Jurafsky. 2021. [Improving factual completeness and consistency of image-to-text radiology report generation](#). In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 5288–5304, Online. Association for Computational Linguistics.
- Maram Mahmoud A Monshi, Josiah Poon, and Vera Chung. 2020. [Deep learning in generating radiology reports: A survey](#). *Artificial Intelligence in Medicine*, 106:101878.
- Michael Moor, Qian Huang, Shirley Wu, Michihiro Yasunaga, Cyril Zakka, Yash Dalmia, Eduardo Pontes Reis, Pranav Rajpurkar, and Jure Leskovec. 2023. [Med-flamingo: a multimodal medical few-shot learner](#). *Preprint*, arXiv:2307.15189.
- Aaron Nicolson, Jason Dowling, and Bevan Koopman. 2023. [Improving chest x-ray report generation by leveraging warm starting](#). *Artificial Intelligence in Medicine*, 144:102633.
- OpenAI, Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman, Diogo Almeida, Janko Altmenschmidt, Sam Altman, Shyamal Anadkat, Red Avila, Igor Babuschkin, Suchir Balaji, Valerie Balcom, Paul Baltescu, Haiming Bao, Mohammad Bavarian, Jeff Belgum, Irwan Bello, Jake Berdine, Gabriel Bernadett-Shapiro, Christopher Berner, Lenny Bogdonoff, Oleg Boiko, Madelaine Boyd, Anna-Luisa Brakman, Greg Brockman, Tim Brooks, Miles Brundage, Kevin Button, Trevor Cai, Rosie Campbell, Andrew Cann, Brittany Carey, Chelsea Carlson, Rory Carmichael, Brooke Chan, Che Chang, Fotis Chantzis, Derek Chen, Sully Chen, Ruby Chen, Jason Chen, Mark Chen, Ben Chess, Chester Cho, Casey Chu, Hyung Won Chung, Dave Cummings, Jeremiah Currier, Yunxing Dai, Cory Decareaux, Thomas Degry, Noah Deutsch, Damien Deville, Arka Dhar, David Dohan, Steve Dowling, Sheila Dunning, Adrien Ecoffet, Atty Eleti, Tyna Eloundou, David Farhi, Liam Fedus, Niko Felix, Simón Posada Fishman, Juston Forte, Isabella Fulford, Leo Gao, Elie Georges, Christian Gibson, Vik Goel, Tarun Gogineni, Gabriel Goh, Rapha Gontijo-Lopes, Jonathan Gordon, Morgan Grafstein, Scott Gray, Ryan Greene, Joshua Gross, Shixiang Shane Gu, Yufei Guo, Chris Hallacy, Jesse Han, Jeff Harris, Yuchen He, Mike Heaton, Johannes Heidecke, Chris Hesse, Alan Hickey, Wade Hickey, Peter Hoeschele, Brandon Houghton, Kenny Hsu, Shengli Hu, Xin Hu, Joost Huizinga, Shantanu Jain, Shawn Jain, Joanne Jang, Angela Jiang, Roger Jiang, Haozhun Jin, Denny Jin, Shino Jomoto, Billie Jonn, Heewoo Jun, Tomer Kaftan, Łukasz Kaiser, Ali Kamali, Ingmar Kanitscheider, Nitish Shirish Keskar, Tabarak Khan, Logan Kilpatrick, Jong Wook Kim, Christina Kim, Yongjik Kim, Jan Hendrik Kirchner, Jamie Kiros, Matt Knight, Daniel Kokotajlo, Łukasz Kondraciuk, Andrew Kondrich, Aris Konstantinidis, Kyle Kosic, Gretchen Krueger, Vishal Kuo, Michael Lampe, Ikai Lan, Teddy Lee, Jan Leike, Jade Leung, Daniel Levy, Chak Ming Li, Rachel Lim, Molly Lin, Stephanie Lin, Mateusz Litwin, Theresa Lopez, Ryan Lowe, Patricia Lue, Anna Makanju, Kim Malfacini, Sam Manning, Todor Markov, Yaniv Markovski, Bianca Martin, Katie Mayer, Andrew Mayne, Bob McGrew, Scott Mayer McKinney, Christine McLeavey, Paul McMillan, Jake McNeil, David Medina, Aalok Mehta, Jacob Menick, Luke Metz, Andrey Mishchenko, Pamela Mishkin, Vinnie Monaco, Evan Morikawa, Daniel Mossing, Tong Mu, Mira Murati, Oleg Murk, David Mély, Ashvin Nair, Reiichiro Nakano, Rameesh Nayak, Arvind Neelakantan, Richard Ngo, Hyeonwoo Noh, Long Ouyang, Cullen O’Keefe, Jakub Pachocki, Alex Paino, Joe Palermo, Ashley Pantuliano, Giambattista Parascandolo, Joel Parish, Emy Parparita, Alex Passos, Mikhail Pavlov, Andrew Peng, Adam Perialman, Filipe de Avila Belbute Peres, Michael Petrov, Henrique Ponde de Oliveira Pinto, Michael, Pokorny, Michelle Pokrass, Vitchyr H. Pong, Tolly Powell, Alethea Power, Boris Power, Elizabeth Proehl, Raul Puri, Alec Radford, Jack Rae, Aditya Ramesh, Cameron Raymond, Francis Real, Kendra Rimbach, Carl Ross, Bob Rotsted, Henri Roussez, Nick Ryder, Mario Saltarelli, Ted Sanders, Shibani Santurkar, Girish Sastry, Heather Schmidt, David Schnurr, John Schulman, Daniel Selsam, Kyla Sheppard, Toki Sherbakov, Jessica Shieh, Sarah Shoker, Pranav

- Shyam, Szymon Sidor, Eric Sigler, Maddie Simens, Jordan Sitkin, Katarina Slama, Ian Sohl, Benjamin Sokolowsky, Yang Song, Natalie Staudacher, Felipe Petroski Such, Natalie Summers, Ilya Sutskever, Jie Tang, Nikolas Tezak, Madeleine B. Thompson, Phil Tillet, Amin Tootoonchian, Elizabeth Tseng, Preston Tuggle, Nick Turley, Jerry Tworek, Juan Felipe Cerón Uribe, Andrea Vallone, Arun Vijayvergiya, Chelsea Voss, Carroll Wainwright, Justin Jay Wang, Alvin Wang, Ben Wang, Jonathan Ward, Jason Wei, CJ Weinmann, Akila Welihinda, Peter Welinder, Jiayi Weng, Lillian Weng, Matt Wiethoff, Dave Willner, Clemens Winter, Samuel Wolrich, Hannah Wong, Lauren Workman, Sherwin Wu, Jeff Wu, Michael Wu, Kai Xiao, Tao Xu, Sarah Yoo, Kevin Yu, Qiming Yuan, Wojciech Zaremba, Rowan Zellers, Chong Zhang, Marvin Zhang, Shengjia Zhao, Tianhao Zheng, Juntang Zhuang, William Zhuk, and Barret Zoph. 2024. [Gpt-4 technical report](#). *Preprint*, arXiv:2303.08774.
- Andreas S. Panayides, Amir Amini, Nenad D. Filipovic, Ashish Sharma, Sotirios A. Tsaftaris, Alistair Young, David Foran, Nhan Do, Spyretta Golemati, Tahsin Kurc, Kun Huang, Konstantina S. Nikita, Ben P. Veasey, Michalis Zervakis, Joel H. Saltz, and Constantinos S. Pattichis. 2020. [AI in medical imaging informatics: Current challenges and future directions](#). *IEEE Journal of Biomedical and Health Informatics*, 24(7):1837–1857.
- Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. 2002. [Bleu: a method for automatic evaluation of machine translation](#). In *Proceedings of the 40th Annual Meeting on Association for Computational Linguistics*, ACL '02, page 311–318, USA. Association for Computational Linguistics.
- Sang-Min Park and Young-Gab Kim. 2023. [Visual language integration: A survey and open challenges](#). *Computer Science Review*, 48:100548.
- Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, Gretchen Krueger, and Ilya Sutskever. 2021. [Learning transferable visual models from natural language supervision](#). *Preprint*, arXiv:2103.00020.
- Samyam Rajbhandari, Jeff Rasley, Olatunji Ruwase, and Yuxiong He. 2020. [Zero: Memory optimizations toward training trillion parameter models](#). *Preprint*, arXiv:1910.02054.
- Andrew B. Sellergren, Christina Chen, Zaid Nabulsi, Yuanzhen Li, Aaron Maschinot, Aaron Sarna, Jenny Huang, Charles Lau, Sreenivasa Raju Kalidindi, Mozziyar Etemadi, Florencia Garcia-Vicente, David Melnick, Yun Liu, Krish Eswaran, Daniel Tse, Neeral Beladia, Dilip Krishnan, and Shravya Shetty. 2022. [Simplified transfer learning for chest radiography models using less data](#). *Radiology*, 305(2):454–465. PMID: 35852426.
- Akshay Smit, Saahil Jain, Pranav Rajpurkar, Anuj Pareek, Andrew Y. Ng, and Matthew P. Lungren. 2020. [Chexbert: Combining automatic labelers and expert annotations for accurate radiology report labeling using bert](#). *Preprint*, arXiv:2004.09167.
- Tim Tanida, Philip Müller, Georgios Kaissis, and Daniel Rueckert. 2023. [Interactive and explainable region-guided radiology report generation](#). In *2023 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. IEEE.
- Rohan Taori, Ishaan Gulrajani, Tianyi Zhang, Yann Dubois, Xuechen Li, Carlos Guestrin, Percy Liang, and Tatsunori B Hashimoto. 2023. [Alpaca: A strong, replicable instruction-following model](#). *Stanford Center for Research on Foundation Models.*, 3(6):7.
- Tao Tu, Shekoofeh Azizi, Danny Driess, Mike Schaeckermann, Mohamed Amin, Pi-Chuan Chang, Andrew Carroll, Chuck Lau, Ryutaro Tanno, Ira Ktena, Basil Mustafa, Aakanksha Chowdhery, Yun Liu, Simon Kornblith, David Fleet, Philip Mansfield, Sushant Prakash, Renee Wong, Sunny Virmani, Christopher Semturs, S Sara Mahdavi, Bradley Green, Ewa Dominowska, Blaise Aguerre y Arcas, Joelle Barral, Dale Webster, Greg S. Corrado, Yossi Matias, Karan Singhal, Pete Florence, Alan Karthikesalingam, and Vivek Natarajan. 2023. [Towards generalist biomedical ai](#). *Preprint*, arXiv:2307.14334.
- Justin Xu, Zhihong Chen, Andrew Johnston, Louis Blankemeier, Maya Varma, Jason Hom, William J. Collins, Ankit Modi, Robert Lloyd, Benjamin Hopkins, Curtis Langlotz, and Jean-Benoit Delbrouck. 2024. Overview of the first shared task on clinical text generation: Rrg24 and “discharge me!”. In *The 23rd Workshop on Biomedical Natural Language Processing and BioNLP Shared Tasks*, Bangkok, Thailand. Association for Computational Linguistics.
- Shawn Xu, Lin Yang, Christopher Kelly, Marcin Sieniek, Timo Kohlberger, Martin Ma, Wei-Hung Weng, Atilla Kiraly, Sahar Kazemzadeh, Zakkai Melamed, Jungyeon Park, Patricia Strachan, Yun Liu, Chuck Lau, Preeti Singh, Christina Chen, Mozziyar Etemadi, Sreenivasa Raju Kalidindi, Yossi Matias, Katherine Chou, Greg S. Corrado, Shravya Shetty, Daniel Tse, Shruthi Prabhakara, Daniel Golden, Rory Pilgrim, Krish Eswaran, and Andrew Sellergren. 2023. [Elixir: Towards a general purpose x-ray artificial intelligence system through alignment of large language models and radiology vision encoders](#). *Preprint*, arXiv:2308.01317.
- Benjamin Yan, Ruochen Liu, David E. Kuo, Subathra Adithan, Eduardo Pontes Reis, Stephen Kwak, Vasantha Kumar Venugopal, Chloe P. O’Connell, Agustina Saenz, Pranav Rajpurkar, and Michael Moor. 2023. [Style-aware radiology report generation with radgraph and few-shot prompting](#). *Preprint*, arXiv:2310.17811.
- Feiyang Yu, Mark Endo, Rayan Krishnan, Ian Pan, Andy Tsai, Eduardo Pontes Reis, Eduardo Kaiser Ururahy Nunes Fonseca, Henrique Min Ho Lee, Zahra Shakeri Hossein Abad, Andrew Y Ng, et al.

2023. Evaluating progress in automatic chest x-ray radiology report generation. *Patterns*, 4(9).

Tianyi Zhang, Varsha Kishore, Felix Wu, Kilian Q. Weinberger, and Yoav Artzi. 2020. [Bertscore: Evaluating text generation with bert](#). *Preprint*, arXiv:1904.09675.