

TiBiX: Leveraging Temporal Information for Bidirectional X-ray and Report Generation

Santosh Sanjeev¹[0000-0003-3664-3844], Fadillah Adamsyah
Maani¹[0000-0001-5927-7782], Arsen Abzhanov¹[0009-0008-4912-7172], Vijay Ram
Papineni²[0000-0002-6162-3290], Ibrahim Almakky¹[0009-0008-8802-7107],
Bartłomiej W. Papież³[0000-0002-8432-2511] and Mohammad
Yaqub¹[0000-0001-6896-1105]

¹ Mohamed bin Zayed University of Artificial Intelligence, Abu Dhabi, UAE

{firstname.lastname}@mbzua.i.ac.ae

² Sheikh Shakhboub Medical City, Abu Dhabi, UAE

{vpapineni}@ssmc.ae

³ Big Data Institute, University of Oxford, UK

{bartlomiej.papiez}@bdi.ox.ac.uk

Abstract. With the emergence of vision language models in the medical imaging domain, numerous studies have focused on two dominant research activities: (1) report generation from Chest X-rays (CXR), and (2) synthetic scan generation from text or reports. Despite some research incorporating multi-view CXRs into the generative process, prior patient scans and reports have been generally disregarded. This can inadvertently lead to the leaving out of important medical information, thus affecting generation quality. To address this, we propose **TiBiX**: Leveraging **T**emporal information for **B**idirectional **X**-ray and Report Generation. Considering previous scans, our approach facilitates bidirectional generation, primarily addressing two challenging problems: (1) generating the current image from the previous image and current report and (2) generating the current report based on both the previous and current images. Moreover, we extract and release a curated temporal benchmark dataset derived from the MIMIC-CXR dataset, which focuses on temporal data. Our comprehensive experiments and ablation studies explore the merits of incorporating prior CXRs and achieve state-of-the-art (SOTA) results on the report generation task. Furthermore, we attain on-par performance with SOTA image generation efforts, thus serving as a new baseline in longitudinal bidirectional CXR-to-report generation. The code is available at <https://github.com/BioMedIA-MBZUAI/TiBiX>.

Keywords: Report Generation · Chest X-ray generation · Multimodal Data · Bidirectional Generation · Longitudinal data

1 Introduction

Image and text generation advancements have propelled the progress in Chest X-ray (CXR) and report generation tasks. The specific focus on CXR stems from

the time-consuming and error-prone nature of the manual report composition task, putting a huge strain on radiologists [12]. Moreover, the quality of manually composed reports depends on the radiologist’s experience, which makes accurate report generation a good aid for inexperienced practitioners. Secondly, to train junior radiologists, high-resolution, diverse X-ray images are needed, but data privacy issues limit access to such images. Developing computer vision methods to generate synthetic, clinically diverse X-rays is highly important. To this effect, some research has been conducted in the domain of CXR image [4,33,24] and report [30,28,23,18] generation.

Transformer-based models like [6,5] have shown promising performance, especially with the introduction of expert tokens in [32]. Strategies such as integrating knowledge graphs [36,11], addressing class-level bias [20], utilizing disease tags [29,27], and applying contrastive learning [19] aim to enhance report quality. In parallel, advancements in natural image generation, driven by techniques like Generative Adversarial Networks (GANs) [9,35], have sparked interest in generating CXR images from radiology reports aiming to improve a model’s interpretability, helping build user trust through visual evidence [16]. However, the translation of radiology reports into CXR images is fraught with challenges, mainly stemming from the linguistic complexities and varied structures of radiology reports, hindering accurate representation in image form [33,4]. Despite efforts to confront these obstacles, existing methods often fail to capture the intricate relationships between textual descriptions and corresponding image features, resulting in suboptimal performance in CXR image generation. Moreover, these methods overlook the temporal aspect of patient data which is very important for medical diagnosis. Temporal CXR generation offers valuable training resources and aids in visual prognosis analysis by utilizing patients’ previous scans.

There are few works exploring the temporal domain for generation tasks [1,37,25]. Unlike the previous works, which focus on only report generation, we focus on bidirectional CXR and report generation. Furthermore, most works require both the prior and current scan images, which is generally not the case in real-life scenarios as patients might not always have prior studies. Recent advancements in bidirectional approaches [15,10] have shown promise in bridging the gap between image-to-text and text-to-image generation tasks. In the medical domain, UniXGen was proposed for bidirectional CXR-to-report generation, leveraging multiple CXR views from the same patient during the same visit to enhance report accuracy [17]. Previous research has focused on generating reports or images at single time points, overlooking the longitudinal aspect of patient data [13].

In this paper, we propose Chest X-ray and Report Bidirectional Generation along the temporal domain. Our main contributions are summarized as follows:

- We propose a novel approach to integrate temporal data for the bidirectional generation task by designing a framework consisting of a transformer-based model equipped with causal attention layers that provide an effective way to learn temporal information while also reducing computational complexity.

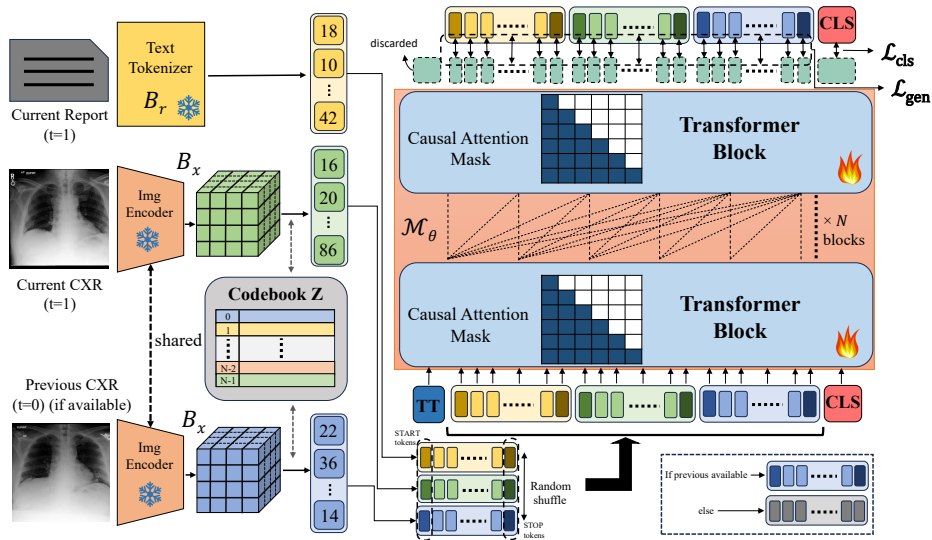


Fig. 1: Our temporal bidirectional CXR-to-report generation framework which can deal with three inputs (i.e. current report, current CXR, and previous CXR). We utilize an image tokenizer B_x and text tokenizer B_r to tokenize CXR image(s) and report, respectively. A transformer-based model \mathcal{M}_θ with *causal* attention is implemented to handle the bidirectional generation task in the *auto-regressive* manner. The input sequence of \mathcal{M}_θ consists of images and text tokens, a temporal token (TT) which encodes the time interval between two consecutive CXR scans, and a learnable *cls* (class) token. We assign TT as the first input sequence and *cls* as the last sequence. During training, the order of the current report, current CXR, and previous CXR are shuffled, while the missing modality is placed at the last part of the input sequence during inference.

We also introduce a *cls* (class) token guidance approach based on the patient’s overall diagnosis and also a temporal token, considering the varying time intervals between patient visits.

- We achieve state-of-the-art (SOTA) performance on the benchmark dataset for the report generation task and achieve comparable performance to the (SOTA) on the image generation task. We also achieve promising results on the classification metrics of images and reports. Furthermore, we support the inclusion of temporal scans by conducting an extensive analysis.
- We extract and release the MIMIC-T dataset, curated from the MIMIC-CXR dataset and consists of longitudinal data of patients extracted from the MIMIC-CXR dataset.

To the best of our knowledge, this is the first effort to include prior(temporal) CXRs and reports for the bidirectional CXR-to-report generation task.

2 Methodology

We formulate the temporal bidirectional CXR-to-report generation task starting from a set of patients $P = \{p_1, p_2, \dots, p_n\}$. For each patient p , we have a set of CXR images $X_p = \{x_{(i)}\}_{i=t_0}^{t_{m_p}}$ along with a set of radiology reports $R_p = \{r_{(i)}\}_{i=t_0}^{t_{m_p}}$, where $m_p \geq 0$ is the total number of time steps and $r_{(t_j)}$ is the corresponding radiology report for $x_{(t_j)}$ captured at time step t_j . Therefore, the total number of samples available is $M = \sum_{p=1}^n (m_p + 1)$. We define the labels for a patient p at time step t_j by $L_{(t_j)} = \{l_{(i)}\}_{i=1}^c$, where $l_{(i)} \in \{0, 1\}$ and c is the total number of possible pathologies. We also define the variable time step $\delta = (t_j - t_{j-1})$ as the time interval between 2 consecutive studies, where $\delta > 0$. As such, the objective of the temporal bidirectional CXR-to-report generation task would be to estimate $r_{(t_j)}$ given $(x_{(t_j)}, x_{t_{(j-1)}})$, while also estimating $x_{(t_j)}$ given $(r_{(t_j)}, x_{t_{(j-1)}})$. To learn this, we employ a transformer-based model with causal attention, which is described later in this section, starting with the multimodal input tokenization and ending with the auto-regressive training process. The high-level visualization of our framework is depicted in Fig. 1.

2.1 Multimodal Tokenization

This work considers three modalities: current report, current CXR, and previous CXR. To facilitate a simple yet effective and memory-efficient multimodal interaction, we tokenize every modality accordingly. We use a CXR image tokenizer $B_x(x) = \mathcal{I}^x \in \mathbb{Z}^{N_x}$ and a report text tokenizer $B_r(r) = \mathcal{I}^r \in \mathbb{Z}^{N_r}$, where N_x and N_r are the number of image and text tokens. We adopt B_x from VQGAN [9] tokenizer, which is composed of two primary components: an image encoder responsible for generating context-rich latent representations and a codebook used to map the representations to a quantized encoding. The main component of our bidirectional generation framework \mathcal{M}_θ is based on a transformer model. We, add unique **START** and **STOP** tokens for every modality, increasing the size of the image and text tokens to $\bar{\mathcal{I}}^x \in \mathbb{Z}^{N_x+2}$ and $\bar{\mathcal{I}}^r \in \mathbb{Z}^{N_r+2}$, respectively to avoid confusion between the modalities, as it would treat all input tokens uniformly. The time interval δ between two consecutive samples is also input to \mathcal{M}_θ as a temporal token (\mathcal{I}^t). This is added to further improve the model’s prognosis capability by mimicking clinicians’ approach to leveraging the time gap between different scans. Finally, we add a learnable token $cls \subset \theta$ to predict the pathologies existing in the patient’s modalities $L_{(t_j)}$, allowing \mathcal{M}_θ to learn from the overall diagnosis.

2.2 Multimodal Generation Framework

Transformer-based Model: \mathcal{M}_θ carries out the temporal bidirectional CXR-to-report generation task as follows:

$$\mathcal{M}_\theta(\mathcal{I}^t, \bar{\mathcal{I}}^{m_1}, \bar{\mathcal{I}}^{m_2}, \bar{\mathcal{I}}^{m_3}) = \hat{\mathcal{I}}^{m_1}, \hat{\mathcal{I}}^{m_2}, \hat{\mathcal{I}}^{m_3}, \hat{L}_t \quad (1)$$

where $\hat{\mathcal{I}}^m \approx \bar{\mathcal{I}}^m$, \hat{L}_t is the prediction of the patient’s diagnosis, and $\{\bar{\mathcal{I}}^{m_1}, \bar{\mathcal{I}}^{m_2}, \bar{\mathcal{I}}^{m_3}\}$ are sampled from set $\{\bar{\mathcal{I}}^{x(t_j)}, \bar{\mathcal{I}}^{r(t_j)}, \bar{\mathcal{I}}^{x(t_{j-1})}\}$ without replacement. Our main training goal is to learn parameters θ so that \mathcal{M}_θ is capable of performing Eq. (1) by optimizing the following objective:

$$\min_{\theta} \frac{1}{M} \sum_{i=1}^M \left(\mathcal{L}_{gen}((\hat{\mathcal{I}}_i^{m_1}, \hat{\mathcal{I}}_i^{m_2}, \hat{\mathcal{I}}_i^{m_3}), (\bar{\mathcal{I}}_i^{m_1}, \bar{\mathcal{I}}_i^{m_2}, \bar{\mathcal{I}}_i^{m_3})) + \lambda \cdot \mathcal{L}_{cls}(\hat{L}_{t_i}, L_{t_i}) \right)$$

where $\mathcal{L}_{gen}(\cdot, \cdot)$ is the generation cross-entropy loss, while $\mathcal{L}_{cls}(\cdot, \cdot)$ is the classification cross-entropy loss based on the *cls*, and $\lambda \in \mathbb{R}$ is a hyperparameter. Some patients do not have a previous scan, i.e. $\bar{\mathcal{I}}^{x(t_{j-1})} = \emptyset$. We replace all missing previous scans with a set of learnable padding tokens to ensure a consistent model input size during training. We follow a similar approach to Eq. (1) during inference without providing the modality we wish to generate, i.e., the input of \mathcal{M}_θ is a patient’s available modalities.

Causal Attention:

To address the computational complexity of the standard attention mechanism, which scales with $\mathcal{O}(N^2)$ for N tokens, we use the Performer architecture along with causal attention to reduce complexity to $\mathcal{O}(N)$. This reduction is achieved by approximating the Softmax function, allowing the multiplication of key and value matrices before multiplying the result by the query matrix. Our proposed model, \mathcal{M}_θ , generates tokens auto-regressively, ensuring each predicted token is used for the next step without future token information leakage. To maintain causality and enable multiprocessing, we implement unidirectional FAVOR+. We add sinusoidal positional embeddings to text embeddings and axial positional embeddings to image embeddings, both previous and current. Additionally, sinusoidal positional embeddings are injected into the full input embeddings before they are passed to the model.

3 Datasets and Experiments

Datasets: We extensively test our method, TiBiX, on a selected subset of MIMIC-CXR [12]. We also introduce MIMIC-T, a new dataset derived from MIMIC-CXR, containing longitudinal patient data. Our preprocessing involves discarding studies lacking both impressions and findings sections and retaining only AP frontal-view scans. This yields approximately 122k, 1k, 2k unique studies from 32521, 255, 266 patients in train, validate, and test splits, respectively. To increase the number of training samples, we utilize prior scans by considering every consecutive pair for the same subject as a new sample. Moreover, we also consider single-study cases simulating real-world scenarios. We prioritize examining previous scans to diagnose diseases and assess prognosis. MIMIC-T provides all previous scans for each patient. The distribution of patients with one or two studies is outlined in Table 1 in Supplementary Material. The dataset, along with the code, is available at <https://github.com/BioMedIA-MBZUAI/TiBiX>.

Implementation Details: We extract impressions and findings from reports

Table 1: Comparison to the previous methods in report generation. (#) - Subset of MIMIC-T with n studies where $n \in [1, 2]$ and $n \in \mathbb{N}$.

Methods	Set (#)	NLG metric						Clinical Efficacy(CE)		
		BLEU-1	BLEU-2	BLEU-3	BLEU-4	ROUGE	METEOR	Precision	Recall	F1
CNN-RNN [21]	CR CX (2)	0.190	0.105	0.066	0.045	0.200	0.096	0.303	0.221	0.230
AdaAtt [22]	CR CX (2)	0.206	0.115	0.072	0.049	0.205	0.102	0.310	0.230	0.234
Transformer [26]	CR CX (2)	0.194	0.109	0.068	0.047	0.200	0.100	0.320	0.228	0.241
M2Transformer[8]	CR CX (2)	0.193	0.108	0.068	0.047	0.200	0.099	<u>0.323</u>	<u>0.248</u>	<u>0.253</u>
R2GenCMN [5]	CR CX (2)	0.348	0.240	<u>0.182</u>	<u>0.146</u>	0.322	<u>0.148</u>	0.321	0.216	0.244
XproNet [31]	CR CX (2)	<u>0.324</u>	0.225	0.172	0.139	<u>0.323</u>	0.141	0.315	0.211	0.240
TiBiX(Ours)	CR CX (1)	0.234	0.156	0.116	0.091	0.272	0.119	0.340	0.190	0.230
	CR CX (2)	0.284	0.184	0.136	0.105	0.267	0.127	0.340	0.220	<u>0.250</u>
	CR (PX,CX) (2)	<u>0.324</u>	<u>0.234</u>	0.185	0.157	0.331	0.162	0.300	<u>0.224</u>	<u>0.25</u>

Table 2: Comparison to state-of-the-art (SOTA) in regard to image generation task. (#) - Subset of MIMIC-T with n studies where $n \in [1, 2]$ and $n \in \mathbb{N}$.

Method	Set (#)	Image Generation metrics				14 diagnosis classification				
		FID	SSIM	MSSIM	IS	Precision	Recall	AUC	Accuracy	F1
Roentgen [4]	GT (2)	-	-	-	-	0.75	0.73	0.79	0.75	0.60
	CX CR (2)	81.02	0.51	0.398	<u>1.96±0.08</u>	0.64	<u>0.64</u>	0.84	0.77	0.60
	Cascaded Diffusion[33]	28.78	0.53	<u>0.428</u>	2.05±0.06	0.48	0.82	<u>0.75</u>	0.66	<u>0.57</u>
TiBiX(Ours)	CX CR (1)	80.48	0.50	0.442	1.70±0.15	0.68	0.28	0.69	0.70	0.42
	CX CR (2)	50.40	0.50	0.438	1.76±0.10	<u>0.61</u>	0.35	0.69	<u>0.74</u>	0.41
	CX (CR,PX) (2)	<u>33.85</u>	<u>0.52</u>	0.472	1.82±0.08	0.56	0.55	<u>0.75</u>	0.77	0.53

and preprocess them by converting text to lowercase, removing punctuation, and digits. We use the CXR-BERT [3] tokenizer with 30,522 unique tokens, including special tokens. AP frontal-view images are resized to 512x512. We employ a pretrained VQ-GAN from [17], trained on MIMIC-CXR with dz=256 and a codebook size of 1024, encoding 512x512 images into 32x32=1024 tokens. Our transformer-based model uses the same configuration of Performer [7]. It handles 1029 image tokens and 30,522 text tokens [9,7]. The model can process two images if the previous one is available. Training lasts for 150 epochs with a batch size of 10, AdamW optimizer, learning rate of 2e-4, weight decay of 1e-2, and cosine decay schedule. Report and image generation use Top-p sampling (p=0.9, temperature=0.7). We ensemble five model weights and select the image with the highest pixel entropy to mitigate stochastic blank image generation. TiBiX can handle multiple prior scans to capture longitudinal effects. Due to computational constraints, we restrict experiments to 1 prior scan, observing good improvement compared to no prior.

4 Results and Discussion

4.1 SOTA Performance Comparison

To demonstrate the effectiveness of our approach, we compare with several SOTA methods in the domain of report and CXR image generation. In all the Tables and Figures, we use the following notations: GT - Ground Truth, CR - Current Report, CX - Current CXR, PX - Previous CXR. (CX|CR) denotes the current CXR generation given the current report, (CX|(CR,PX)) denotes the current

CXR generation given the current report and previous CXR.

Report Generation: We compare our results on the report generation task with several SOTA image-captioning works such as the CNN-RNN[21], Transformer [26], AdaAtt [22], M2Transformer [8] in the natural domain as well as the SOTA works in the medical domain including R2GenCMN [5] and XproNet [31]. For [34,21,26,8] methods, we rerun them on our datasets, following the official implementation of [8]. We adhere to the same preprocessing steps as in our approach, while keeping the other hyperparameters consistent with those specified in their codebase. Table 1 shows that compared to other works, our method achieves SOTA results on most of the NLG metrics and also achieves an on-par performance on the classification scores.

CXR Image Generation: Few works in the medical domain have focused on the generation of Chest X-ray images. We compare our results with the recent top works in the field of X-ray generation, which have used Diffusion Models(DM) [4,33]. We have implemented these pretrained models on the subset of the test set having 2 studies, and the results are shown in Table 2. RoentGen is a medical domain-adapted latent diffusion model based on the Stable Diffusion pipeline whereas Cascaded Diffusion(CD) [33] improves upon this work, by having cascaded models. The first model in [33] generates a lower resolution image and the second one generates a higher resolution image, enhancing the first DM’s output. While TiBiX shows significant improvement in generation metrics, against [4] pretrained weights (also evident from images in Supplementary), classification metrics are lower. There are works exploring issues with [4] like [14] that points out data leakage from test into training data & [2] also points out issues regarding validity, fairness and bias. Roentgen’s [4] high classification scores, surpassing those of real images, may imply that pathology recognition is easy & lacks the diversity & complexity of real-world cases. Compared to CD [33] we see similar results in almost all cases except Recall. This is due to an increase in failure cases in imbalanced classes (+ to - ratio) such as Pneumothorax & Lung Lesion. However, this is not the case for CD’s [33] VAE & super-resolution components as they were trained on a dataset 3.5x larger than MIMIC-T.

It is worth noting that all the works we compare are specifically designed for a single task, either reporting or image generation. In contrast, our framework excels in bidirectional generation while achieving state-of-the-art performance across all tasks. Tab. 1 & Tab. 2 also shows generated reports and images w/o prior scans, demonstrating model’s generalizability. We also provide ablation study in Tables 4 and 5 in Supplementary Material.

4.2 Self-Comparison

We analyze subsets of the MIMIC-T dataset to assess the impact of including a prior CXR scan on report generation. In Table 1, we evaluate report quality and clinical efficacy (CE) on two test subsets: (1) patients with 1 study and (2) patients with 2 studies. In the second case, we analyze CR | CX and CR | (PX, CX). Including a prior scan with the current report increases BLEU-4 by about 1.5x and improves all metrics, including CE. This aligns with the observation

that radiologists produce more accurate reports with multiple scans. In Table 2, we compare image generation quality across MIMIC-T subsets using FID, SSIM, MSSIM, and IS metrics. We also use a pretrained DenseNet121 classifier on the images to get classification scores. The generation quality improves significantly with both prior CXR and CR, achieving the best results in all evaluation metrics, including classification scores.




Fig. 2: (A) Report Generation Task: The first case describes an example of (1) current report generation (CR|CX) (2) current report generation (CR|(PX,CX)). (B) CXR Image Generation task: The left side of the report is (CX|CR) and the on the right side is (CX|(CR,PX)). The GT is present for comparison.

4.3 Qualitative Analysis

The benefits of considering the prior scan when generating a report or a scan are demonstrated in Fig. 2. Fig. 2-A illustrates that in the case where the model does not consider the prior scan it produces a report with false findings that are not in the ground truth (e.g. the report describes cardiomegaly when it's not there). Furthermore, it hallucinates and includes the word "unchanged" even though the previous scan is not provided. On the other hand, when the model considers the previous scan it generates a report with findings more aligned with the ground truth (highlighted in different colors). Fig. 2-B illustrates that when the model considers the previous scan, it generates a scan with an anatomical structure more similar to the ground truth. The resultant scan is also closer to the ground truth in terms of portraying the pathologies. For example, "the opacification of the right lung base as opposed to the left lung being clear is more accurate when the models take into account the previous scan". When the

models do not have the prior scan, the resultant image deviates from the ground truth more. We also provide extensive qualitative analysis comparing previous methods in Tables 2 and 3 in the Supplementary Material.

5 Conclusion

In this work, we introduce a novel approach to enhance image and report generation by integrating temporal data from previous CXR scans. Our framework includes a transformer-based model with causal attention, facilitating bidirectional CXR-to-report generation while considering the time difference between consecutive scans. Experimental results on the MIMIC-T dataset, a temporal subset of MIMIC-CXR, demonstrate promising performance. Qualitative analysis and ablation studies justify the inclusion of prior scans and highlight the importance of each architectural component. While our work lays a foundational framework, future enhancements could explore incorporating prior reports and domain expertise through knowledge graphs. Furthermore, we aim to create a metric capable of capturing the temporal aspect, addressing the inadequacies of current metrics in this regard.

Disclosure of Interests. The authors have no competing interests to declare that are relevant to the content of this article.

References

1. Bannur, S., Hyland, et al.: Learning to exploit temporal structure for biomedical vision-language processing. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 15016–15027 (2023)
2. Bhardwaj, et al.: On the notion of hallucinations from the lens of bias and validity in synthetic cxr images. arXiv preprint arXiv:2312.06979 (2023)
3. Boecking, B., Usuyama, et al.: Making the most of text semantics to improve biomedical vision-language processing. In: European conference on computer vision. pp. 1–21. Springer (2022)
4. Chambon, et al.: Roentgen: vision-language foundation model for chest x-ray generation. arXiv preprint arXiv:2211.12737 (2022)
5. Chen, Z., Shen, et al.: Cross-modal memory networks for radiology report generation. arXiv preprint arXiv:2204.13258 (2022)
6. Chen, Z., Song, et al.: Generating radiology reports via memory-driven transformer. arXiv preprint arXiv:2010.16056 (2020)
7. Choromanski, et al.: Rethinking attention with performers. arXiv preprint arXiv:2009.14794 (2020)
8. Cornia, et al.: Meshed-memory transformer for image captioning. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 10578–10587 (2020)
9. Esser, et al.: Taming transformers for high-resolution image synthesis. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 12873–12883 (2021)

10. Huang, et al.: Du-vlg: Unifying vision-and-language generation via dual sequence-to-sequence pre-training. arXiv preprint arXiv:2203.09052 (2022)
11. Huang, et al.: Kiut: Knowledge-injected u-transformer for radiology report generation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 19809–19818 (2023)
12. Johnson, et al.: MIMIC-CXR, a de-identified publicly available database of chest radiographs with free-text reports. *Sci. Data* **6**(317), 1–8 (Dec 2019)
13. Kayser, et al.: Explaining Chest X-Ray Pathologies in Natural Language. In: MICCAI 2022: 25th International Conference, Singapore, September 18–22, 2022, Proceedings, Part V, pp. 701–713. Springer-Verlag, Berlin, Germany (Sep 2022)
14. Khosravi, et al.: Synthetically enhanced: unveiling synthetic data’s potential in medical imaging research. *EBioMedicine* **104** (2024)
15. Kim, et al.: L-verse: Bidirectional generation between image and text. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 16526–16536 (2022)
16. Lanfredi, et al.: Adversarial Regression Training for Visualizing the Progression of Chronic Obstructive Pulmonary Disease with Chest X-Rays. In: MICCAI 2019 (2019)
17. Lee, H., Kim, et al.: Unified chest x-ray and radiology report generation model with multi-view chest x-rays. arXiv preprint arXiv:2302.12172 (2023)
18. Li, et al.: A self-guided framework for radiology report generation. In: Wang, et al. (eds.) MICCAI. Springer Nature Switzerland (2022)
19. Liu, et al.: Contrastive attention for automatic chest X-ray report generation. In: Zong, et al. (eds.) Findings of the Association for Computational Linguistics: ACL-IJCNLP 2021. pp. 269–280. Association for Computational Linguistics, Online (Aug 2021)
20. Liu, F., Wu, X., Ge, S., Fan, W., Zou, Y.: Exploring and distilling posterior and prior knowledge for radiology report generation. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 13753–13762 (2021)
21. Liu, G., Hsu, et al.: Clinically accurate chest x-ray report generation. In: Machine Learning for Healthcare Conference. pp. 249–269. PMLR (2019)
22. Lu, J., Xiong, et al.: Knowing when to look: Adaptive attention via a visual sentinel for image captioning. arxiv 2016. arXiv preprint arXiv:1612.01887 (2016)
23. Müller, et al.: Radiological reports improve pre-training for localized imaging tasks on chest x-rays. In: MICCAI 2022. pp. 647–657. Springer Nature Switzerland, Cham (2022)
24. Packhäuser, K., et al.: Generation of anonymous chest radiographs using latent diffusion models for training thoracic abnormality classification systems. In: 2023 IEEE 20th International Symposium on Biomedical Imaging (ISBI). pp. 1–5. IEEE (2023)
25. Serra, D., et al.: Controllable chest X-ray report generation from longitudinal representations. In: Bouamor, H., Pino, J., Bali, K. (eds.) Findings of the Association for Computational Linguistics: EMNLP 2023. pp. 4891–4904. Association for Computational Linguistics, Singapore (Dec 2023)
26. Vaswani, et al.: Attention is all you need **30** (2017)
27. Wang, et al.: A self-boosting framework for automated radiographic report generation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 2433–2442 (2021)
28. Wang, et al.: An inclusive task-aware framework for radiology report generation. In: Wang, et al. (eds.) MICCAI. pp. 568–577. Springer Nature Switzerland, Cham (2022)

29. Wang, et al.: A medical semantic-assisted transformer for radiographic report generation. In: MICCAI. pp. 655–664. Springer (2022)
30. Wang, et al.: A medical semantic-assisted transformer for radiographic report generation. In: Wang, et al. (eds.) MICCAI. Springer Nature Switzerland (2022)
31. Wang, J., Bhalerao, A., He, Y.: Cross-modal prototype driven network for radiology report generation. In: European Conference on Computer Vision. pp. 563–579. Springer (2022)
32. Wang, Z., Liu, L., Wang, L., Zhou, L.: Metransformer: Radiology report generation by transformer with multiple learnable expert tokens. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 11558–11567 (2023)
33. Weber, et al.: Cascaded latent diffusion models for high-resolution chest x-ray synthesis. In: Pacific-Asia Conference on Knowledge Discovery and Data Mining. pp. 180–191. Springer (2023)
34. Xu, K., Ba, et al.: Show, attend and tell: Neural image caption generation with visual attention. In: International conference on machine learning. pp. 2048–2057. PMLR (2015)
35. Xu, T., et al.: AttnGAN: Fine-grained text to image generation with attentional generative adversarial networks. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 1316–1324 (2018)
36. Yang, et al.: Knowledge matters: Chest radiology report generation with general and specific knowledge. *Medical image analysis* **80**, 102510 (2022)
37. Zhu, et al.: Utilizing longitudinal chest x-rays and reports to pre-fill radiology reports. In: MICCAI 2023. pp. 189–198. Springer Nature Switzerland, Cham (2023)