



This MICCAI paper is the Open Access version, provided by the MICCAI Society. It is identical to the accepted version, except for the format and this watermark; the final published version is available on SpringerLink.

Advancing Text-Driven Chest X-Ray Generation with Policy-Based Reinforcement Learning

Woojung Han^{*1} Chanyoung Kim^{*2} Dayun Ju¹
Yumin Shim³ Seong Jae Hwang^{†2}

¹ Department of Computer Science, Yonsei University, Seoul, Republic of Korea

² Department of Artificial Intelligence, Yonsei University, Seoul, Republic of Korea

³ College of Medicine, Yonsei University, Seoul, Republic of Korea

{dnwjddl, chanyoung, juda0707, symin0425, seongjae}@yonsei.ac.kr

Abstract. Recent advances in text-conditioned image generation diffusion models have begun paving the way for new opportunities in the modern medical domain, in particular, particularly in generating Chest X-rays (CXRs) from diagnostic reports. Nonetheless, to further drive the diffusion models to generate CXRs that faithfully reflect the complexity and diversity of real data, it has become evident that a nontrivial learning approach is needed. In light of this, we propose CXRL, a framework motivated by the potential of reinforcement learning (RL). Specifically, we integrate a policy gradient RL approach with well-designed multiple distinctive CXR-domain specific reward models. This approach guides the diffusion denoising trajectory, achieving precise CXR posture and pathological details. Here, considering the complex medical image environment, we present “RL with Comparative Feedback” (RLCF) for the reward mechanism, a human-like comparative evaluation that is known to be more effective and reliable in complex scenarios compared to direct evaluation. Our CXRL framework includes jointly optimizing learnable adaptive condition embeddings (ACE) and the image generator, enabling the model to produce more accurate and higher perceptual CXR quality. Our extensive evaluation of the MIMIC-CXR-JPG dataset demonstrates the effectiveness of our RL-based tuning approach. Consequently, our CXRL generates pathologically realistic CXRs, establishing a new standard for generating CXRs with high fidelity to real-world clinical scenarios. Project page: <https://micv-yonsei.github.io/cxrl2024/>

Keywords: Chest X-Ray · Diffusion Models · Reinforcement Learning

1 Introduction

Diffusion probabilistic models (DPMs) [9, 28] have revolutionized conditional image synthesis [26], particularly in the realm of medical imaging [5, 16, 24, 25],

^{*} Equal contribution [†] Corresponding author

[‡] In accordance with the MIMIC-CXR data usage license [13], the text reports presented in Fig.1, 2 and 3 have been rephrased while maintaining the original content.

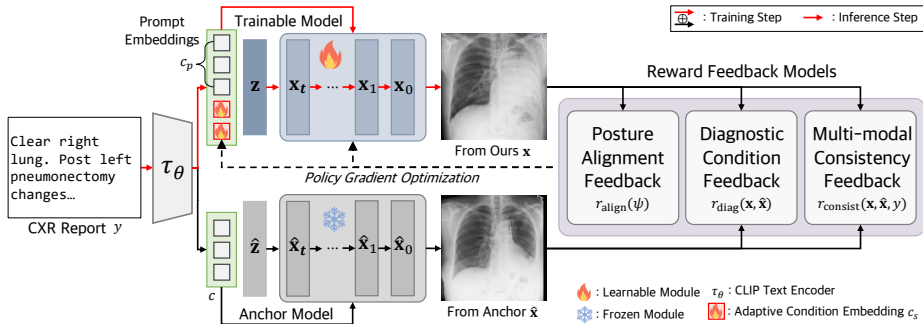


Fig. 1. The overview of **CXRL**. Our model employs policy gradient optimization with multi-reward feedback to fine-tune both the image generator and ACE, producing realistic and accurate CXRs that closely match the input report.[‡]

by enabling text-driven generation [3, 12]. This innovative approach addresses medical field challenges, such as data scarcity and privacy concerns, offering a promising avenue for synthetic medical data utilization in various applications (i.e., privacy-free medical image data for medical education or research) [14].

However, generating medical images via text conditions is challenging with diffusion models, mainly due to their inability to accurately represent subtle diagnostic differences, while medical images have complex characteristics (i.e., minor tissue texture variations per patient or disease) [14]. This issue emerges because of the reliance on approximating the log-likelihood objective of the training data, prompting the model to predominantly generate images that align with the distribution of the training data [2, 6], thus failing to accurately represent minor variations characteristic of medical images. Meanwhile, recent computer vision studies suggest treating diffusion denoising as a multi-step decision-making problem [2, 6, 7, 18], allowing the use of policy gradient-based reinforcement learning (RL) which can overcome the existing limitations. Motivated by such potential, we first integrate policy-based RL into medical image synthesis by tuning an image generator to enhance text-to-medical image generation using non-differentiable reward functions. However, our empirical observations indicated that solely optimizing the image generator limits the scope of pathological expressions. Consequently, we introduce **CXRL**, jointly fine-tuning the image generator and *learnable adaptive condition embeddings* (ACE) to condition the image generation process flexibly, enriching pathological features represented.

Although prior works [2, 6, 7, 18] have successfully incorporated data-driven human feedback models [15, 17] into RL for image generation, applying human feedback models to medical image generation task presents unique and significant challenges. These challenges primarily stem from the “complexity of the medical imaging environment”, which makes it difficult for humans to quantitatively express the overall quality of medical images with absolute numbers [22, 27]. Given these challenges, particularly in the context of complex and uncertain

environments like medical imaging, the findings of Mussweiler and Posten [23] stand out, illustrating that humans are better at making consistent and accurate judgments through “relative comparison” rather than absolute evaluation.

Therefore, we adopt *Reinforcement Learning with “Comparative” Feedback* (RLCF) method, similar to [7], which enhances the reliability of the feedback mechanism by human-like comparative evaluation processes, a suitable feedback for complex text-to-medical image generation environments (Fig. 1). Specifically, RLCF assesses image pairs generated from ours and the anchor model, providing rewards for surpassing the anchor’s images and penalties when our model’s images fall short of the anchor’s, which we categorize as negative actions. This strategy employs both rewards and penalties, steering our model away from producing images inferior to the anchor’s images. To sum up, our RLCF framework not only mimics human feedback for a reliable training experience but also integrates penalties to deter negative actions in synthesizing medical images.

Focusing specifically on Chest X-ray (CXR) generation, this paper addresses three key challenges in generating report-conditioned CXRs tackling each with specific feedback to achieve our final goal. **(i)** Posture Alignment Feedback: We ensure accurate posture alignment through a specialized network, addressing potential misalignments that could impair disease detection and diagnostics, such as missing crucial areas (i.e., costophrenic angle). **(ii)** Diagnostic Condition Feedback: To enhance diagnostic accuracy in generated CXRs, we use feedback from a state-of-the-art pretrained CXR lesion classifier [4] which guides the model to generate more precise pathologic features. **(iii)** Multimodal Consistency Feedback: To ensure that CXRs generated from reports are semantically in agreement with the reports, we provide feedback using large-scale representations [29] trained on CXR-report pairs. This feedback guides our CXR generation model to overcome disagreement between generated CXRs and reports.

Contributions. **(1)** Our study pioneers in applying RL to text-conditioned medical image synthesis, particularly in CXRs, focusing on detail refinement and input condition control for clinical accuracy. **(2)** We advance report-to-CXR generation with a RLCF-based rewarding mechanism, emphasizing posture alignment, pathology accuracy, and consistency between input reports and generated CXRs. **(3)** We jointly optimize the image generator and ACE via reward feedback models, ensuring image-text alignment and medical accuracy across varied reports, setting a new benchmark in a report-to-CXR generation.

2 Methods

In our study, we integrate policy gradient RL with diffusion models to enhance the denoising process. We first describe the preliminaries and detail our framework (Fig. 1) guided by a carefully designed multi-reward system (Fig. 2).

2.1 Preliminaries

Conditional Diffusion Probabilistic Models. Conditional Diffusion Probabilistic Models (Conditional DPMs) [26] improve upon DPMs [9, 28] by inte-




Fig. 2. A detailed illustration of our reward feedback models. We incorporate three different feedbacks for report-to-CXR generation model to generate goal-oriented CXRs.[‡]

grating conditions like text prompts into the reverse diffusion process, guiding the synthesis towards high diversity and adherence to specified features. These models train on a conditional log-likelihood to adjust outputs closely to these conditions, enhancing diversity and targeted representation capability.

RL-based Diffusion Model Fine-tuning. The RL paradigm interprets *the diffusion denoising process as a series of actions* aiming to maximize a reward function that evaluates the adherence of generated images to a given prompt [2]. The sequential decision-making (denoising process) is formulated by an objective function that targets the cumulative reward across the diffusion trajectory. Given the context of RL-based fine-tuning of conditional DPMS, the objective function \mathcal{J}_θ is defined as $\mathcal{J}_\theta = \mathbb{E}_{p(\mathbf{y})} \mathbb{E}_{p_\theta(\mathbf{x}|\mathbf{y})} [r(\mathbf{x}, \mathbf{y})]$, where $p(\mathbf{y})$ is the distribution of the text condition \mathbf{y} , $p_\theta(\mathbf{x}|\mathbf{y})$ is the distribution produced by the pretrained DPM conditioned on text \mathbf{y} , and $r(\mathbf{x}, \mathbf{y})$ is the reward for image \mathbf{x} given text condition \mathbf{y} . As in DPOK [6], the policy gradient of \mathcal{J}_θ w.r.t. the parameters θ is

$$\nabla_\theta \mathcal{J}_\theta = \mathbb{E} \left[r(\mathbf{x}, \mathbf{y}) \sum_{t=1}^T \nabla_\theta \log p_\theta(\mathbf{x}_{t-1} | \mathbf{x}_t, \mathbf{y}, t) \right], \quad (1)$$

where T is the total number of time steps in the diffusion trajectory, \mathbf{x}_t is the state of the image at time step t , and the gradient ∇_θ is taken w.r.t. the model parameters θ . By applying policy gradients, DPMS iteratively fine-tune their parameters, improving image quality and enabling the generation of images with high fidelity, detail, and diversity.

2.2 Reinforcement Learning with Comparison Feedback

Unlike existing methods that rely on direct feedback [2, 6, 18], we employ Reinforcement Learning with Comparative Feedback (RLCF), contrasting model performance with a static anchor model, as illustrated in Fig. 1 and Fig. 2. Utilizing RLCF, our model adapts through positive feedback for outperforming,

and negative for underperformance, compared to the frozen anchor. With given reward feedback signals, we jointly optimize two learning aspects:

(I) Image Generator: The image generator (diffusion model) aims to create the final CXR \mathbf{x} through a series of denoising steps, conditioned by the text prompt embeddings c extracted from CXR report \mathbf{y} . We fine-tune the diffusion layers in a RL manner, advancing their fidelity in medical imaging.

(II) Learnable Adaptive Condition Embeddings: Given the variability in radiologist reporting styles, relying solely on a report itself for CXR generation presents significant challenges. For this reason, inspired by Lester *et al.* [20], we combine original static prompt embeddings $c_p = \tau_\theta(\mathbf{y}) \in \mathbb{R}^{M \times d_\tau}$ with the *learnable adaptive condition embeddings* (ACE), represented as $c_s \in \mathbb{R}^{N \times d_\tau}$. It is achieved by concatenating them to form the input condition as $c = [c_s, c_p] \in \mathbb{R}^{(N+M) \times d_\tau}$. Here, τ_θ symbolizes the parameterized frozen text encoder for conditioning image generator, \mathbf{y} is the input report, d_τ is the embedding dimension, N is a sequence length of c_s and M denotes sequence length of c_p . ACE and c_p jointly allow the image generator more flexible inputs via cross-attention, enabling a precise capture of diagnostic and posture alignment features in CXRs.

2.3 Goal-oriented Reward Feedback Models

Our reward system comprises three models with distinct objectives (Fig. 2).

(I) Posture Alignment Feedback. Generated CXRs often face scaling issues, such as excessive zooming or rotation, obscuring essential details. To counter these undesirable effects, we introduce a reward signal to align the CXR’s posture with a canonical orientation to preserve essential parts. The posture alignment model is trained using L2 loss between the real CXRs and canonical CXR, where canonical CXR is created by averaging 500 random CXRs from the dataset, following Liu *et al.* [21]. The posture alignment model adjusts CXRs by predicting parameters $\psi = \{s_x, s_y, t_x, t_y, \Theta\}$ for an affine transformation, including scaling (s_x, s_y) , translation (t_x, t_y) , and rotation Θ . The posture alignment reward, $r_{\text{align}}(\psi)$, quantifies the magnitude of transformation obtained from the frozen posture alignment model. Our r_{align} is calculated in independent parameter space as $r_{\text{align}}(\psi) = -\{(\max(|s_x - 1|, |s_y - 1|) + \frac{|\Theta|}{2\pi} + \sqrt{t_x^2 + t_y^2})\}$. It aligns CXRs to the canonical pose, enhancing posture precision and visibility of essential regions.

(II) Diagnostic Condition Feedback. To accurately reflect generated CXRs with referenced pathologies, we classify them using parsed report label, rewarding its accuracy. To keep away from negative actions, we integrate the RLCF, comparing our model’s CXR \mathbf{x} against anchor model’s CXR $\hat{\mathbf{x}}$, using their score difference for reward. Our diagnostic condition feedback r_{diag} is defined by the formula: $r_{\text{diag}}(\mathbf{x}, \hat{\mathbf{x}}) = \{\text{accuracy}(\mathcal{G}(\mathbf{x})) - \text{accuracy}(\mathcal{G}(\hat{\mathbf{x}}))\}$, where \mathcal{G} represents a state-of-the-art pretrained CXR lesion classifier network [4].

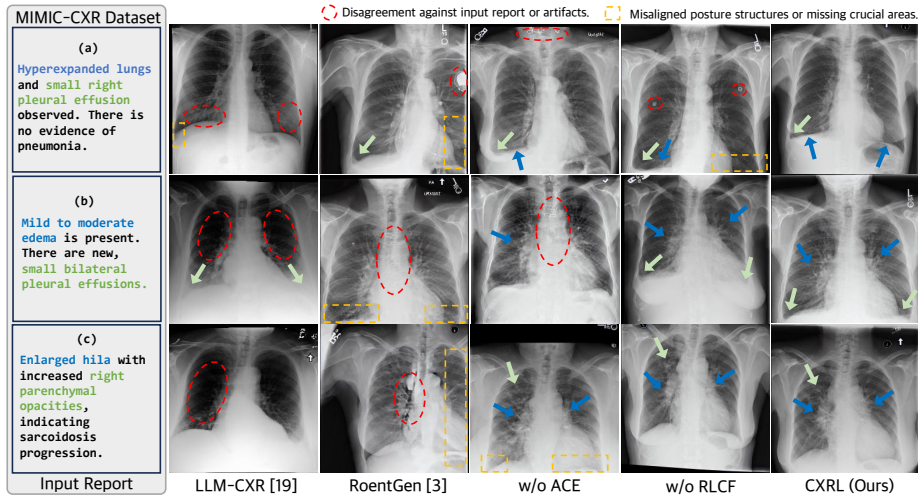


Fig. 3. Comparison between previous state-of-the-art report-to-CXR generation models [19, 3] and ours. The blue and green texts match their corresponding colored arrows.[‡]

(III) Multimodal Consistency Feedback. Finally, we enforce the generated CXRs to better match their reports. We leverage a multimodal latent representation [29] pretrained with CXR-report pairs for semantic agreement assessment. Specifically, we employ RLCF to compute the difference in similarity scores between generated CXR \mathbf{x} by our model and the report \mathbf{y} , and the generated CXR $\hat{\mathbf{x}}$ by an anchor model and the same report. The reward feedback is defined as $r_{\text{consist}}(\mathbf{x}, \hat{\mathbf{x}}, \mathbf{y}) = \{d_{\text{cosine}}(\mathcal{F}_{\text{visual}}(\mathbf{x}), \mathcal{F}_{\text{text}}(\mathbf{y})) - d_{\text{cosine}}(\mathcal{F}_{\text{visual}}(\hat{\mathbf{x}}), \mathcal{F}_{\text{text}}(\mathbf{y}))\}$, where d_{cosine} represents the cosine similarity and $\mathcal{F}_{\text{visual}}$ and $\mathcal{F}_{\text{text}}$ represent the visual and text encoders, respectively, mapping feature embeddings from their respective modal inputs to joint multimodal latent space [29].

2.4 Optimizing Policy Gradient

Integrating the introduced reward feedbacks, we update the policy gradient in a RL manner, by expanding Eq. (1). Our final reward to maximize is $r(\mathbf{x}, \hat{\mathbf{x}}, \mathbf{y}, \psi) = \lambda_{\text{align}} r_{\text{align}}(\psi) + \lambda_{\text{diag}} r_{\text{diag}}(\mathbf{x}, \hat{\mathbf{x}}) + \lambda_{\text{consist}} r_{\text{consist}}(\mathbf{x}, \hat{\mathbf{x}}, \mathbf{y})$, where λ_{align} , λ_{diag} , and λ_{consist} are hyperparameters. With given rewards, for B images in mini-batches and T total diffusion denoising steps, we update the gradient as follows:

$$\nabla \mathcal{J}_{\theta} = \frac{1}{B} \sum_{i=1}^B \sum_{t=1}^T (r(\mathbf{x}^{(i)}, \hat{\mathbf{x}}^{(i)}, \mathbf{y}^{(i)}, \psi^{(i)})) \times \nabla_{\theta} \log p_{\theta}(\mathbf{x}_{t-1}^{(i)} | \mathbf{x}_t^{(i)}, c^{(i)}, t). \quad (2)$$

Table 1. Evaluation of generated CXRs from multiple feedback perspectives. The table compares the performance of various methods using three evaluation metrics.

Method	(a) Posture Alignment $r_{\text{align}}(\psi)$ (\uparrow)	(b) Classification AUROC (\uparrow)	(c) Semantic Consistency CXR-CLIP Distance (\uparrow)
RoentGen [3]	-0.363	0.65	0.285
LLM-CXR [19]	-0.327	<u>0.78</u>	0.201
CXRL (Ours)	-0.246	0.81	0.343
w/o ACE	-0.354	0.66	<u>0.321</u>
w/o RLCF (direct feedback)	<u>-0.319</u>	0.71	0.301

3 Experiments

Datasets. We utilize PA and AP views from the MIMIC-CXR-JPG dataset [13], while limiting report impressions to 74 tokens or less. We randomly sample 5,000 pairs from the official train split for training. For testing, similarly to the training set, we utilize PA views from the test split, resulting in 1,000 pairs.

Implementation Details. CXRL fine-tunes with LoRA layers [10] on RoentGen-initialized Stable Diffusion [3, 26], whereas the anchor model uses static RoentGen. We train a posture alignment model with the MIMIC-CXR-JPG dataset [13] using ResNet-18 [8]. A DenseNet-121 [11] CXR classifier [4] provides diagnostic feedback. $\mathcal{F}_{\text{text}}$ and $\mathcal{F}_{\text{vision}}$ utilize CXR-CLIP [29], with BioClinicalBERT [1] and ResNet50 [8] serving as the backbone, respectively. DDPO algorithm [2] guides policy gradient optimization, with 81 CXRs per batch, learning rate of $3\text{e-}4$. λ_{align} , λ_{diag} , and λ_{consist} are at 1, 10, and 10, with $N = 3$ and $M \leq 74$.

3.1 Evaluation Results

Qualitative Analysis. In Fig. 3, we qualitatively compare our method to previous models [3, 19]. **(1) *Disagreement against input report or artifacts (red circle):*** Our CXRL exhibits superior agreement between CXR and report compared to the baselines, especially in their capability to avoid generating unmentioned objects. For instance, in Fig. 3(a), while RoentGen deviates from the report content by generating a pacemaker, CXRL maintains semantic consistency by not doing so. Furthermore, in case Fig. 3(b), RoentGen generates sternal wires (red circle) that are not mentioned in the report, whereas CXRL does not. In terms of diagnosing specific conditions, CXRL shows a high level of agreement. For example, in case Fig. 3(c), while LLM-CXR and RoentGen either fail to identify or incorrectly generate hilar enlargement (blue colored text) with artifacts, CXRL accurately represents them (blue arrow). **(2) *Misaligned posture structures or missing crucial areas (yellow box):*** While others occasionally miss details, our CXRL consistently shows the thoracic anatomy comprehensively, including lung apices and costophrenic angles. In contrast, RoentGen often loses critical areas due to inadequate posture alignment.

Quantitative Evaluation on CXR Specific Quality. Table 1 shows the evaluation of our generated CXR across three metrics, reflecting the rewards

Table 2. Comparative analysis of generated CXR quality: (a) compares established models using FID and MS-SSIM metrics; (b) evaluates the impact of reward components on FID scores.

(a) Quality Assessment			(b) Ablation: Reward Feedback Models		
Model	FID (\downarrow)	MS-SSIM (\downarrow)	Exp. #	Reward	FID (\downarrow)
RoentGen [3]	2.58	0.197	-	Anchor Model [3]	2.58
LLM-CXR [19]	4.01	0.275	1	+ r_{align}	1.89
CXRL (Ours)	1.51	0.147	2	+ r_{diag}	1.64
w/o ACE	1.79	0.192	3	+ r_{consist}	1.92
w/o RLCF	<u>1.58</u>	<u>0.188</u>	4	$r_{\text{align}} + r_{\text{diag}} + r_{\text{consist}}$	1.51

we targeted, which are key qualities of CXR. (a) The r_{align} measures posture alignment, where values closer to 0 indicate a better fit. (b) The AUROC score assesses diagnostic accuracy, and (c) the semantic consistency score evaluates the alignment between CXR and the report. Table 1 confirms our rewards are enhancing report-to-CXR generation model towards the goal-oriented result.

Quantitative Evaluation on Image Quality. Table 2a details the fidelity of our generated CXRs, using the FID and MS-SSIM metrics for benchmarking against baselines. Employing a CXR-specific DenseNet-121 [4] for FID, CXRL outperforms baseline models [3, 19], showcasing superior image realism. Our MS-SSIM scores, indicative of structural accuracy, also prove competitive, affirming our model’s performance in generating authentic CXR textures.

Medical Expert Assessment. In pursuit of medical expert feedback, we assessed 20 generated CXRs with five radiology experts. For impartiality, our CXRL was fairly compared with baselines [19, 3] on the same hardware/software configuration, and the generated CXRs were rated against reports on a 1 to 5 scale. Average scores were: RoentGen **3.4**, LLM-CXR **2.4**, and CXRL at **4.1**, showing superior CXR-report alignment and realism of our generated CXRs.

3.2 Ablation Study

Effect of ACE and RLCF. We validate the effectiveness of ACE (Adaptive Condition Embeddings) and RLCF (comparative reward feedback) by excluding the corresponding factor from fully combined CXRL. Table 1 presents a comparison of key CXR quality metrics in our CXRL framework with and without the integration of ACE and RLCF. While removing either ACE or RLCF improves performance over the anchor model [3], ours with both ACE and RLCF achieves the best overall performance. Similar aspects are shown in terms of image quality assessment (Table 2(a)). In particular, w/o RLCF, which uses direct feedback optimization with ACE, seems to improve CXR quality compared to the anchor by providing flexible input conditions. This experiment demonstrates that both ACE and RLCF are essential contributions of our model, in both domain-specific quality and image quality aspects (See Fig. 3 for qualitative result).

Effect of Reward Models. In Table 2(b), we detail how each reward signal influences model performance by examining FID scores. Individual rewards en-

hance CXR quality (Exp. #1 to #3) beyond anchor levels which we aim to surpass. Combining all rewards (Exp. #4) synergistically lowers the FID, showing that multiple feedback signals crucially enhance CXR synthesis by capturing broader features for high-fidelity CXRs (See supplement for qualitative result).

4 Conclusion and Discussion

In this study, we present **CXRL** for report-to-CXR generation using RLCF, which significantly enhances medical image synthesis through a policy gradient RL and an advanced adaptable reward feedback system. Our exhaustive experiments demonstrate our framework’s state-of-the-art performance in generating CXRs. The proposed adaptable reward framework not only excels in CXR generation but may also pioneer a wide spectrum of applications in various domains of medical image synthesis, offering enhanced diagnostics, contributing to medical education and research, and ensuring ethical access to medical imaging data.

Disclosure of Interests. The authors have no competing interests.

Acknowledgments. We thank S. Jung, J.E. Lee, S.H. Yun, and C. Lee for their valuable medical expertise and advice. This work was supported in part by the IITP 2020-0-01361 (AI Graduate School Program at Yonsei University), NRF RS-2024-00345806, and NRF RS-2023-00219019 funded by Korean Government (MSIT).

References

1. Alsentzer, E., Murphy, J., Boag, W., Weng, W.H., Jin, D., Naumann, T., McDermott, M.: Publicly available clinical BERT embeddings. In: Proceedings of the 2nd Clinical Natural Language Processing Workshop. pp. 72–78. Association for Computational Linguistics, Minneapolis, Minnesota, USA (Jun 2019)
2. Black, K., Janner, M., Du, Y., Kostrikov, I., Levine, S.: Training diffusion models with reinforcement learning. In: The Twelfth International Conference on Learning Representations (2024)
3. Chambon, P., Bluethgen, C., Delbrouck, J.B., Van der Sluijs, R., Połacin, M., Chaves, J.M.Z., Abraham, T.M., Purohit, S., Langlotz, C.P., Chaudhari, A.: Roentgen: vision-language foundation model for chest x-ray generation. arXiv preprint arXiv:2211.12737 (2022)
4. Cohen, J.P., Viviano, J.D., Bertin, P., Morrison, P., Torabian, P., Guarrera, M., Lungren, M.P., Chaudhari, A., Brooks, R., Hashir, M., et al.: Torchxrayvision: A library of chest x-ray datasets and models. In: International Conference on Medical Imaging with Deep Learning. pp. 231–249. PMLR (2022)
5. Du, Y., Jiang, Y., Tan, S., Wu, X., Dou, Q., Li, Z., Li, G., Wan, X.: Arsdm: colonoscopy images synthesis with adaptive refinement semantic diffusion models. In: International conference on medical image computing and computer-assisted intervention. pp. 339–349. Springer (2023)

6. Fan, Y., Watkins, O., Du, Y., Liu, H., Ryu, M., Boutilier, C., Abbeel, P., Ghavamzadeh, M., Lee, K., Lee, K.: Reinforcement learning for fine-tuning text-to-image diffusion models. In: Thirty-seventh Conference on Neural Information Processing Systems (2023)
7. Hao, Y., Chi, Z., Dong, L., Wei, F.: Optimizing prompts for text-to-image generation. In: Thirty-seventh Conference on Neural Information Processing Systems (2023)
8. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 770–778 (2016)
9. Ho, J., Jain, A., Abbeel, P.: Denoising diffusion probabilistic models. *Advances in neural information processing systems* **33**, 6840–6851 (2020)
10. Hu, E.J., yelong shen, Wallis, P., Allen-Zhu, Z., Li, Y., Wang, S., Wang, L., Chen, W.: LoRA: Low-rank adaptation of large language models. In: International Conference on Learning Representations (2022)
11. Huang, G., Liu, Z., Van Der Maaten, L., Weinberger, K.Q.: Densely connected convolutional networks. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 4700–4708 (2017)
12. Jiang, L., Mao, Y., Wang, X., Chen, X., Li, C.: Cola-diff: Conditional latent diffusion model for multi-modal mri synthesis. In: International Conference on Medical Image Computing and Computer-Assisted Intervention. pp. 398–408. Springer (2023)
13. Johnson, A.E., Pollard, T.J., Greenbaum, N.R., Lungren, M.P., Deng, C.y., Peng, Y., Lu, Z., Mark, R.G., Berkowitz, S.J., Horng, S.: Mimic-cxr-jpg, a large publicly available database of labeled chest radiographs. *arXiv preprint arXiv:1901.07042* (2019)
14. Kazerouni, A., Aghdam, E.K., Heidari, M., Azad, R., Fayyaz, M., Ilker: Diffusion models in medical imaging: A comprehensive survey. *Medical Image Analysis* **88**, 102846 (2023)
15. Ke, J., Ye, K., Yu, J., Wu, Y., Milanfar, P., Yang, F.: Vila: Learning image aesthetics from user comments with vision-language pretraining. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 10041–10051 (2023)
16. Khader, F., Mueller-Franzes, G., Arasteh, S.T., Han, T., Haarbuerger, C., Schulze-Hagen, M., Schad, P., Engelhardt, S., Baessler, B., Foersch, S., et al.: Medical diffusion–denoising diffusion probabilistic models for 3d medical image generation. *arXiv preprint arXiv:2211.03364* (2022)
17. Kirstain, Y., Polyak, A., Singer, U., Matiana, S., Penna, J., Levy, O.: Pick-a-pic: An open dataset of user preferences for text-to-image generation. In: Thirty-seventh Conference on Neural Information Processing Systems (2023)
18. Lee, S.H., Li, Y., Ke, J., Yoo, I., Zhang, H., Yu, J., Wang, Q., Deng, F., Entis, G., He, J., et al.: Parrot: Pareto-optimal multi-reward reinforcement learning framework for text-to-image generation. In: European Conference on Computer Vision. Springer (2024)
19. Lee, S., Kim, W.J., Chang, J., Ye, J.C.: LLM-CXR: Instruction-finetuned LLM for CXR image understanding and generation. In: The Twelfth International Conference on Learning Representations (2024)
20. Lester, B., Al-Rfou, R., Constant, N.: The power of scale for parameter-efficient prompt tuning. In: Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing. pp. 3045–3059. Association for Computational Linguistics, Online and Punta Cana, Dominican Republic (Nov 2021)

21. Liu, J., Zhao, G., Fei, Y., Zhang, M., Wang, Y., Yu, Y.: Align, attend and locate: Chest x-ray diagnosis via contrast induced attention network with limited supervision. In: Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV) (October 2019)
22. Margaret Cheng, H.L., Stikov, N., Ghugre, N.R., Wright, G.A.: Practical medical applications of quantitative mr relaxometry. *Journal of Magnetic Resonance Imaging* **36**(4), 805–824 (2012)
23. Mussweiler, T., Posten, A.C.: Relatively certain! comparative thinking reduces uncertainty. *Cognition* **122**(2), 236–240 (2012)
24. Peng, W., Adeli, E., Zhao, Q., Pohl, K.M.: Generating realistic 3d brain mris using a conditional diffusion probabilistic model. In: International conference on medical image computing and computer-assisted intervention. Springer (2023)
25. Pinaya, W.H., Tudosiu, P.D., Dafflon, J., Da Costa, P.F., Fernandez, V., Nachev, P., Ourselin, S., Cardoso, M.J.: Brain imaging generation with latent diffusion models. In: MICCAI Workshop on Deep Generative Models. pp. 117–126. Springer (2022)
26. Rombach, R., Blattmann, A., Lorenz, D., Esser, P., Ommer, B.: High-resolution image synthesis with latent diffusion models. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 10684–10695 (2022)
27. Rosen, A.F., Roalf, D.R., Ruparel, K., Blake, J., Seelaus, K., Villa, L.P., Ciric, R., Cook, P.A., Davatzikos, C., Elliott, M.A., et al.: Quantitative assessment of structural image quality. *Neuroimage* **169**, 407–418 (2018)
28. Song, J., Meng, C., Ermon, S.: Denoising diffusion implicit models. In: International Conference on Learning Representations (2021)
29. You, K., Gu, J., Ham, J., Park, B., Kim, J., Hong, E.K., Baek, W., Roh, B.: Cxr-clip: Toward large scale chest x-ray language-image pre-training. In: International Conference on Medical Image Computing and Computer-Assisted Intervention. pp. 101–111. Springer (2023)