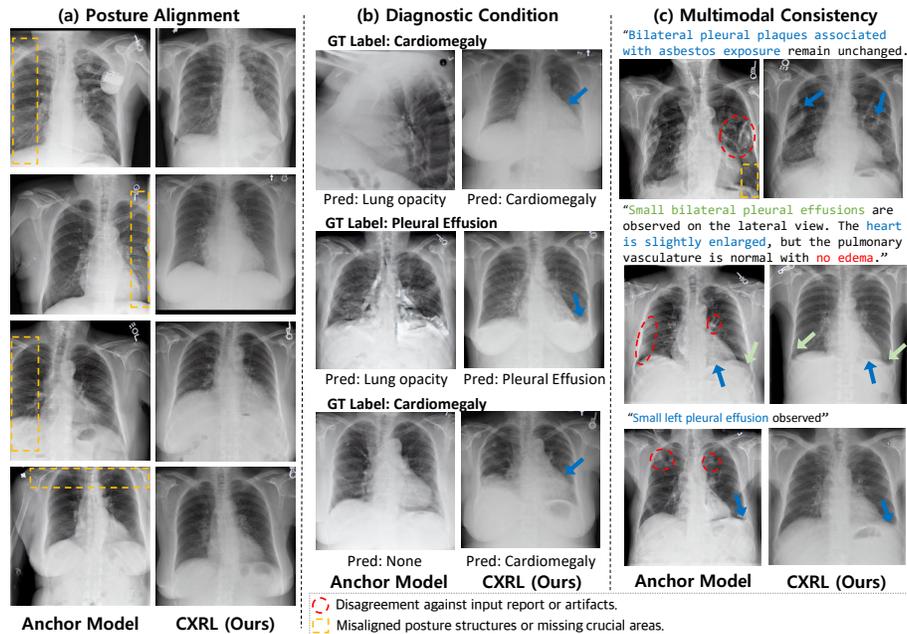


# Supplementary Materials, Advancing Text-Driven Chest X-Ray Generation With Policy-Based Reinforcement Learning

Woojung Han\* Chanyoung Kim\* Dayun Ju  
Yumin Shim Seong Jae Hwang†



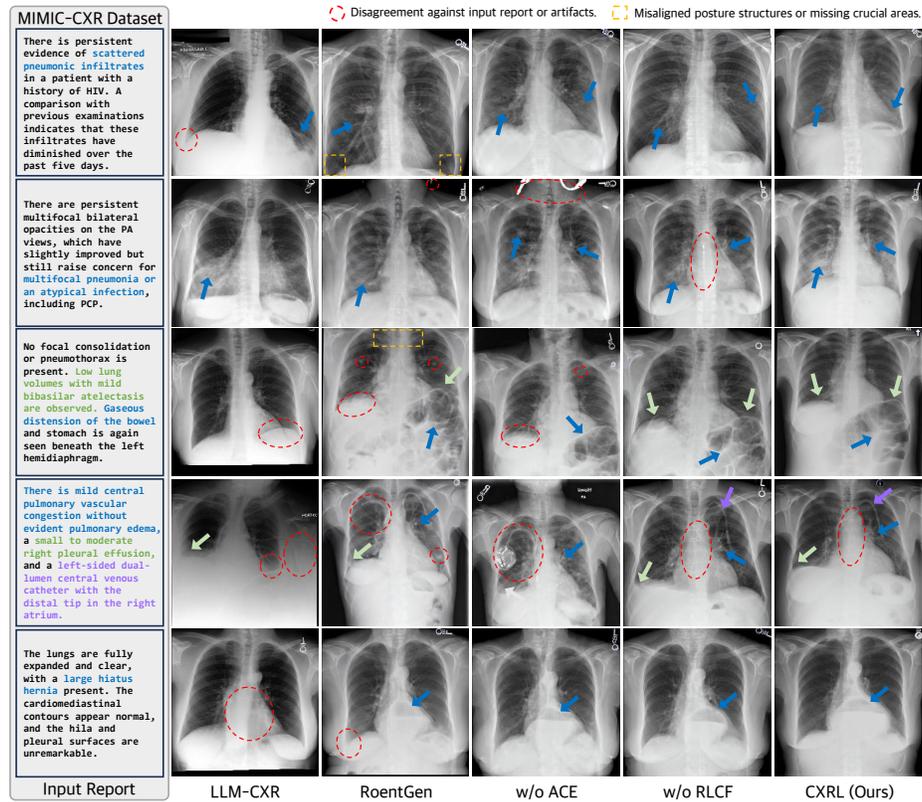
**Fig. 1.** Qualitative ablation on each reward models. (a): CXRL shows significantly better alignment of the *clavicle* and *costophrenic angle* compared to anchor regarding posture alignment. (b): CXRL demonstrates improved predictive diagnostic accuracy, closely matching the GT and enhancing clinical decision-making. (c): The multimodal consistency reward ensures that CXRs and reports correspond well, as observed by arrows and text in matching colors. Conversely, anchor sometimes generates content in the report that is either irrelevant or should not be present. Specifically, in the second row, the presence of edema is verifiable within the red circle, indicating misalignment between the CXR and the report.<sup>‡</sup>

\* Equal contribution    † Corresponding author

<sup>‡</sup> In accordance with the MIMIC-CXR data usage license, the text reports presented in Fig.1 and 2 have been rephrased while maintaining the original content.

**Table 1.** AUROC scores for each pathology class, obtained by processing generated CXRs from various models through a state-of-the-art CXR pathology classifier. Notably, our **CXRL** model outperforms other baselines and even surpasses the diagnostic accuracy of original real CXRs (GT) in the dataset, highlighting the proficiency of our approach in capturing the details of each pathology accurately.

AUROC ( $\uparrow$ )	Atel.	Cnsl.	Pmtx.	Edem.	Eff.	Pna.	Cmgl.	Les.	Frac.	Opac.	ECm.	Avg.
RoentGen	.78	.90	.79	.71	.46	.39	.46	.57	.90	.48	.71	.65
LLM-CXR	.86	.88	.81	.86	.68	<b>.76</b>	<b>.67</b>	.76	.88	.62	.82	.78
<b>CXRL</b>	<b>.87</b>	<b>.94</b>	<b>.88</b>	<b>.89</b>	<b>.77</b>	.60	.63	<b>.79</b>	<b>.94</b>	<b>.70</b>	<b>.86</b>	<b>.81</b>



**Fig. 2.** Additional qualitative results of our framework comparing against baselines. The colored texts match their corresponding colored arrows. Ours w/o ACE or RLCF demonstrates superior report agreement and posture alignment compared to other baselines, and **CXRL** is observed to generate more advanced high-fidelity CXRs that highlight our methodology’s effectiveness in synthesizing clinically accurate medical images.<sup>‡</sup>