

Supplemental Material

1 Implementation details

For all training purposes, 1 RTX 2080 GPU workstation is used. For both Baseline-LO and LO+direction, L-BFGS optimizer is used with learning rate of 0.001. $W+$ vectors corresponding to layers 2,3 and 4 are optimized for opacity and vectors corresponding to layers 3,4 and 5 are optimized for cardiomegaly, values of c and σ are kept as 0.25 and 0.30 respectively, $\alpha = 2.0, \beta = 0.3, \gamma = 0.1, \theta' = 35^\circ$. All optimizations run for 100 epochs. We run classification experiments with a pretrained Resnet50 model for 15 epochs with Adam optimizer. Learning rate is set to 0.01 stepwise decreasing every 5 epochs by a factor of 10. Binary Cross Entropy metric is utilized for all training.

2 More Results From The Proposed Method

Go to the next page.




Fig. 1: Direction used is to decrease opacity and increase cardiac size i.e PC10




Fig. 2: Direction used is to decrease opacity and decrease cardiac size i.e PC12




Fig. 3: Direction used is to increase opacity and increase cardiac size i.e PC20




Fig. 4: Direction used is to increase opacity and increase cardiac size i.e PC20

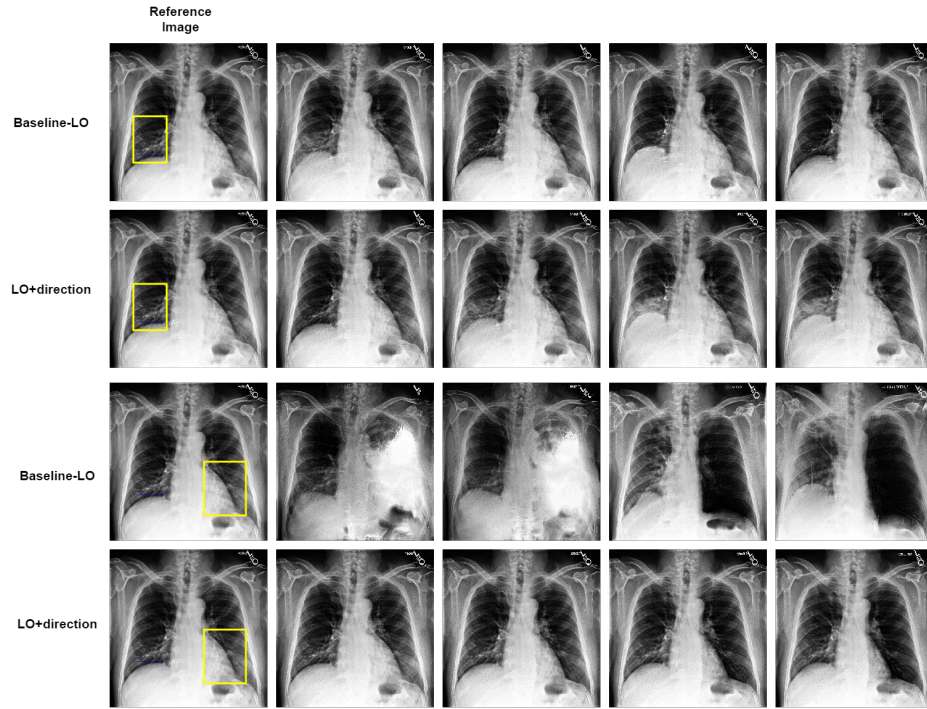


Fig. 5: Direction used in proposed method is to increase opacity and decrease cardiac size i.e PC15

3 Clinical Evaluation

We present a sample questionnaire provided to the pulmonologist for the evaluation of the generated images by our proposed method. Some artifacts are present in the images such as small aberrations and minor lines across the lung lobe. While most of the images are consistent with the disease of cardiomegaly and opacity, only one image is identified to look unrealistic with respect to opacity and one with below average image quality(image rating: 2). Go to the next page.




IMAGE 1:
Quality of the chest X-ray(on a scale of 1-5) : **4**
Opacity present: yes/no: **NO**
Cardiomegaly present: Yes/no: **YES**
Comments (optional):: -



IMAGE 2:
Quality of the chest X-ray(on a scale of 1-5) : **4**
Opacity present: yes/no: **NO**
Cardiomegaly present: Yes/no: **YES**
Comments (optional):: -




IMAGE 1:
Quality of the chest X-ray(on a scale of 1-5): **3**
Opacity present: yes/no: **YES**
Cardiomegaly present: Yes/no: **NO**
Comments (optional): **Artifacts on the Left side**



IMAGE 2:
Quality of the chest X-ray(on a scale of 1-5) : **4**
Opacity present: yes/no: **YES**
Cardiomegaly present: Yes/no: **YES**
Comments (optional):: -

Fig. 6: sample questionnaire with images and their corresponding response



IMAGE 1:
 Quality of the chest X-ray(on a scale of 1-5): **3**
 Opacity present: yes/no: **NO**
 Cardiomegaly present: Yes/no: **YES**
 Comments (optional): -



IMAGE 2:
 Quality of the chest X-ray(on a scale of 1-5) : **4**
 Opacity present: yes/no: **YES**
 Cardiomegaly present: Yes/no: **NO**
 Comments (optional): **Opacity does not feel real**



IMAGE 1:
 Quality of the chest X-ray(on a scale of 1-5): **3**
 Opacity present: yes/no: **NO**
 Cardiomegaly present: Yes/no: **NO**
 Comments (optional): -



IMAGE 2:
 Quality of the chest X-ray(on a scale of 1-5) : **3**
 Opacity present: yes/no: **NO**
 Cardiomegaly present: Yes/no: **YES**
 Comments (optional): -

Fig. 7: sample questionnaire with images and their correponding response

4 Generated Images

We present a set of randomly sampled generated images from our trained Style-GAN. The images posits that the model does not suffer from mode collapse and thus, it is suitable for our experiments. Go to next page



Fig. 8: Randomly sampled generated image

Table 1: Quantification of variation in the sequences using PV metric for Cardiomegaly and Opacity.

	$\alpha = 0$	$\beta = 0$	$\gamma = 0$	$\alpha = 2.0, \beta = 0.3, \gamma = 0.1$
Cardiomegaly	0.08 ± 0.11	-0.16 ± 0.18	0.32 ± 0.21	0.51 ± 0.20
Opacity	0.05 ± 0.12	-0.20 ± 0.16	0.28 ± 0.13	0.44 ± 0.14

Table 2: Quantification of variation in the sequences using PV metric for different values of θ' .

	$\theta' = 35^\circ$	$\theta' = 53^\circ$	$\theta' = 84^\circ$
Cardiomegaly	0.51 ± 0.20	0.38 ± 0.12	0.11 ± 0.16
Opacity	0.44 ± 0.14	0.22 ± 0.13	0.06 ± 0.20

5 Ablation Studies

PV metric has been defined to quantify the monotonicity and localised changes in the sequence generated. Since the hyper-parameters $\alpha, \beta, \gamma, \theta'$ directly affects the aforementioned attributes of the sequence, we ablate them w.r.t PV metric and give example UMAPs and images where necessary. Unless stated otherwise, assume $\alpha = 2.0, \beta = 0.3, \gamma = 0.1, \theta' = 35^\circ, \sigma = 0.3$. In Table 1, we present PV metric values by assigning one of the hyper-parameter equal to 0. Each value represents the mean of the PV metric along with its standard deviation, obtained from 50 different normal images subjected to the latent optimization to generate 12 semantically edited images. We can infer that the maximum PV value is achieved when none of the hyper-parameters are equal to zero (col. 4). For *eg.* when α is set to zero (\mathcal{L}_X is eliminated), we see very negligible changes across the sequence (Fig. 9(b)), when β is set to zero ($\mathcal{L}_{\text{spring}}(Z; 1)$ is eliminated), we see optimization oscillating between two similar images (Fig. 9(c)) and when γ is set to zero ($\mathcal{L}_{\text{spring}}(Z; 2)$ is eliminated), we are able to achieve monotonicity to some extent but its PV value is less than the PV value when all the parts of the loss are used (Fig. 9(a) and Fig. 9(d)).

We also investigate different values of θ' and present their corresponding PV values in Table 2. As we increase the value of θ' , PV value decreases because the proposed LO+direction becomes more similar to Baseline LO. Best results are achieved with $\theta' = 0.3$. Fig. 10(b) presents UMAPs for varied values of θ' for an example image.

We also present UMAPs for different values of σ for the same example image in Fig. 10(a). As we increase the value of σ , we observe decrease in PV value. We believe that images with required changes in the Bounding Box lie close to the original image. As we increase σ , we force the optimization to go farther from the original point in the direction of u , but due to unavailability of the required points, it has to return towards the original point hindering the monotonicity of the sequence and hence, a low PV value.




Fig. 9: UMAP and sequence generated using LO+Direction for different values of hyper-parameters. Values between two consecutive points in the UMAP represent distance between them. Please zoom in to see them.

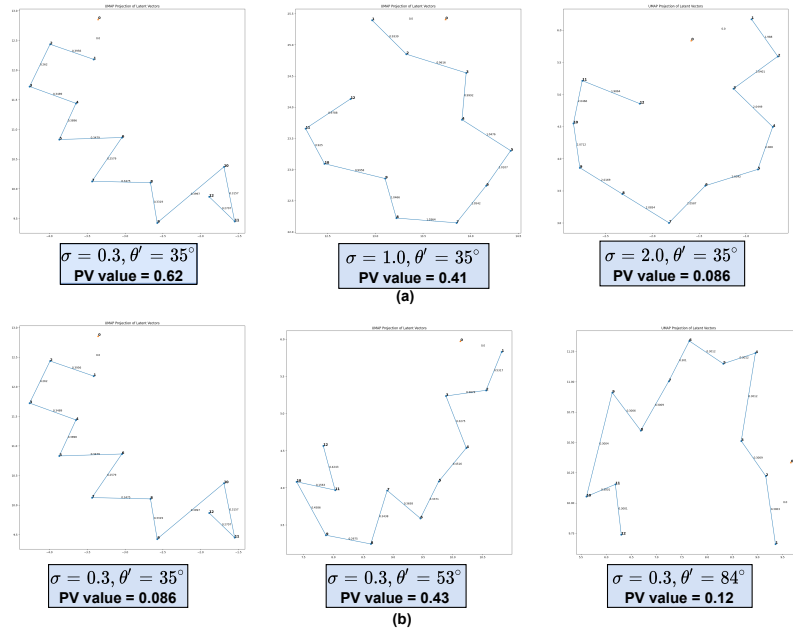


Fig. 10: UMAP for different values of σ and θ' . Values between two consecutive points in the UMAP represent distance between them. Please zoom in to see them.




Fig. 11: A failure case of the latent optimization where no significant changes are seen across the sequence.

6 Limitation

In some cases, ill-positioned latent points may hamper the latent optimization. Fig. 11 shows one such instance where no significant changes are observed in the generated sequence with any combination of hyper-parameters.