# Supplementary Material
## Ablation Study and Registration Result

## Ablation Study

This supplementary material provides extended ablation experiments evaluating the design choices of AortaDiff. All experiments use the same dataset splits and evaluation metrics as described in the main manuscript except explicitely stated.

### Ablation variants

We evaluate the following design choices:

- **Reconstruction Loss**: use MSE or perceptual loss for CECT image reconstruction.
- **Mask representation**: use signed distance function to represent lumen mask instead of binary mask.
- **Multi-task loss combination**: use the uncertainty-aware multi-task loss instead of using the equal-weighted sum of segmentation and generation loss.

From the reconstruction results in Table 1, we observe that using LPIPS loss yields better reconstruction performance compared to MSE loss across all three models. This improvement is likely due to the background misalignments between our NCCT–CECT pairs, where a pixel-wise loss forces the model to learn these misalignments instead of focusing on perceptually meaningful features. For both multi-task models, AortaDiff-F and AortaDiff-P, the results in Table 2 show that using an SDF-based mask representation yields better segmentation performance compared to a binary mask. In addition, Table 3 shows that using the uncertainty-aware loss leads to improved performance for both tasks compared to the equal-weight loss formulation.

Table 1: Ablation results comparing different reconstruction losses, using either MSE or perceptual loss.

| Reconstruction Loss | PSNR (dB) ↑ | SSIM ↑ | LPIPS ↓ |
|---|---|---|---|
| CDM with MSE loss | 20.87 | 0.7253 | 0.1077 |
| CDM with LPIPS loss | 23.80 | 0.7616 | 0.0859 |
| CycleGAN with MSE loss | 22.83 | 0.6492 | 0.1283 |
| CycleGAN with LPIPS loss | 23.19 | 0.6535 | 0.1129 |
| U-Net with MSE loss | 23.83 | 0.7602 | 0.0971 |
| U-Net with LPIPS loss | 24.40 | 0.7785 | 0.0815 |

Table 2: Ablation results for different loss combination strategies.

| Mask Representation | PSNR (dB) ↑ | SSIM ↑ | LPIPS ↓ | Lumen Dice ↑ |
|---|---|---|---|---|
| AortaDiff-F Binary | 25.41 | 0.8262 | 0.0688 | 0.8753 |
| AortaDiff-F SDF | 25.61 | 0.8385 | 0.0671 | 0.8887 |
| AortaDiff-P Binary | 25.47 | 0.8198 | 0.0661 | 0.8760 |
| AortaDiff-P SDF | 25.48 | 0.8296 | 0.0606 | 0.8933 |

Table 3: Ablation results for multi-task loss. All metrics are the same as in the main paper. PSNR, SSIM, and LPIPS are computed over the aorta region.

| Multi-task Loss | PSNR (dB) ↑ | SSIM ↑ | LPIPS ↓ | Lumen Dice ↑ |
| --- | --- | --- | --- | --- |
| AortaDiff-F with Equal-Weight Loss | 25.37 | 0.8207 | 0.0660 | 0.8700 |
| AortaDiff-F with Uncertainty-Aware Loss | 25.61 | 0.8385 | 0.0671 | 0.8887 |
| AortaDiff-P with Equal-Weight Loss | 24.40 | 0.7808 | 0.0757 | 0.8830 |
| AortaDiff-P with Uncertainty-Aware Loss | 25.48 | 0.8296 | 0.0606 | 0.8933 |

# SyN-Based Non-Rigid Registration Result

This section explains the reason that we did not apply non-rigid registration as a preprocessing step. Our task focuses on accurately translating the aortic region. However, after applying non-rigid registration (SyN in ANTs (Advanced Normalization Tools)), the aorta segmentation Dice score between the registered non-contrast CT and the original contrast CT decreases, despite the overall images are better aligned (Table 4). This indicates that the registration process improves global alignment but fails to keep the alignment of the aortic region. Therefore, non-rigid registration is not suitable for our application. We address the misalignment issue by introducing the aorta-focused generator, as described in the main paper.

Table 4: OxAAA registration results before and after image registration. PSNR, SSIM, and LPIPS are computed over the whole image.

| **Registration** | **Aorta Dice ↑** | **PSNR ↑** | **SSIM ↑** | **LPIPS ↓** |
| --- | --- | --- | --- | --- |
| Before Registration | 0.95 | 20.34 | 0.64 | 0.11 |
| After Registration | 0.84 | 28.67 | 0.86 | 0.059 |