Synthesis of Spectral CT Tissue Maps from Single-Energy CT Using Conditional Diffusion Models

Jerry Sarubbe
Department of Electrical Engineering, Stanford University

Motivation

- Spectral CT enables material-specific tissue decomposition, allowing clinicians to isolate tissue types that would otherwise be indistinguishable on conventional CT.
- Clinically, this information is valuable but requires expensive dual-energy imaging systems and exposes patients to higher doses of radiation.
- We propose using a Conditional DDPM (C-DDPM) to synthesize 3-channel tissue maps (adipose, fibroglandular, calcification) from a single low-energy (50 kVp) CT scan, eliminating the need for dual-energy acquisition.
- Dataset: AAPM DL-Spectral-CT Challenge – paired breast tissue phantom slices (512×512), split 85/10/5 for train/val/test. [1]

Related Work

- Material decomposition CNNs have estimated dual-energy CT from low-energy CT, demonstrating that learned mappings can be derived from standard CT acquisitions, however, these approaches can produce blurry outputs, especially for finer structures. [2]
- Denoising Diffusion Probabilistic Models generate high-fidelity images through iterative denoising enabling practical use in medical imaging. [3]
- Conditional DDPM applications for spectral CT were demonstrated to be viable for various applications such as generating contrast-enhanced dual-energy CT images from non-contrast single-energy CT as well as for iodine map synthesis – outperforming existing methods. [4]

Method

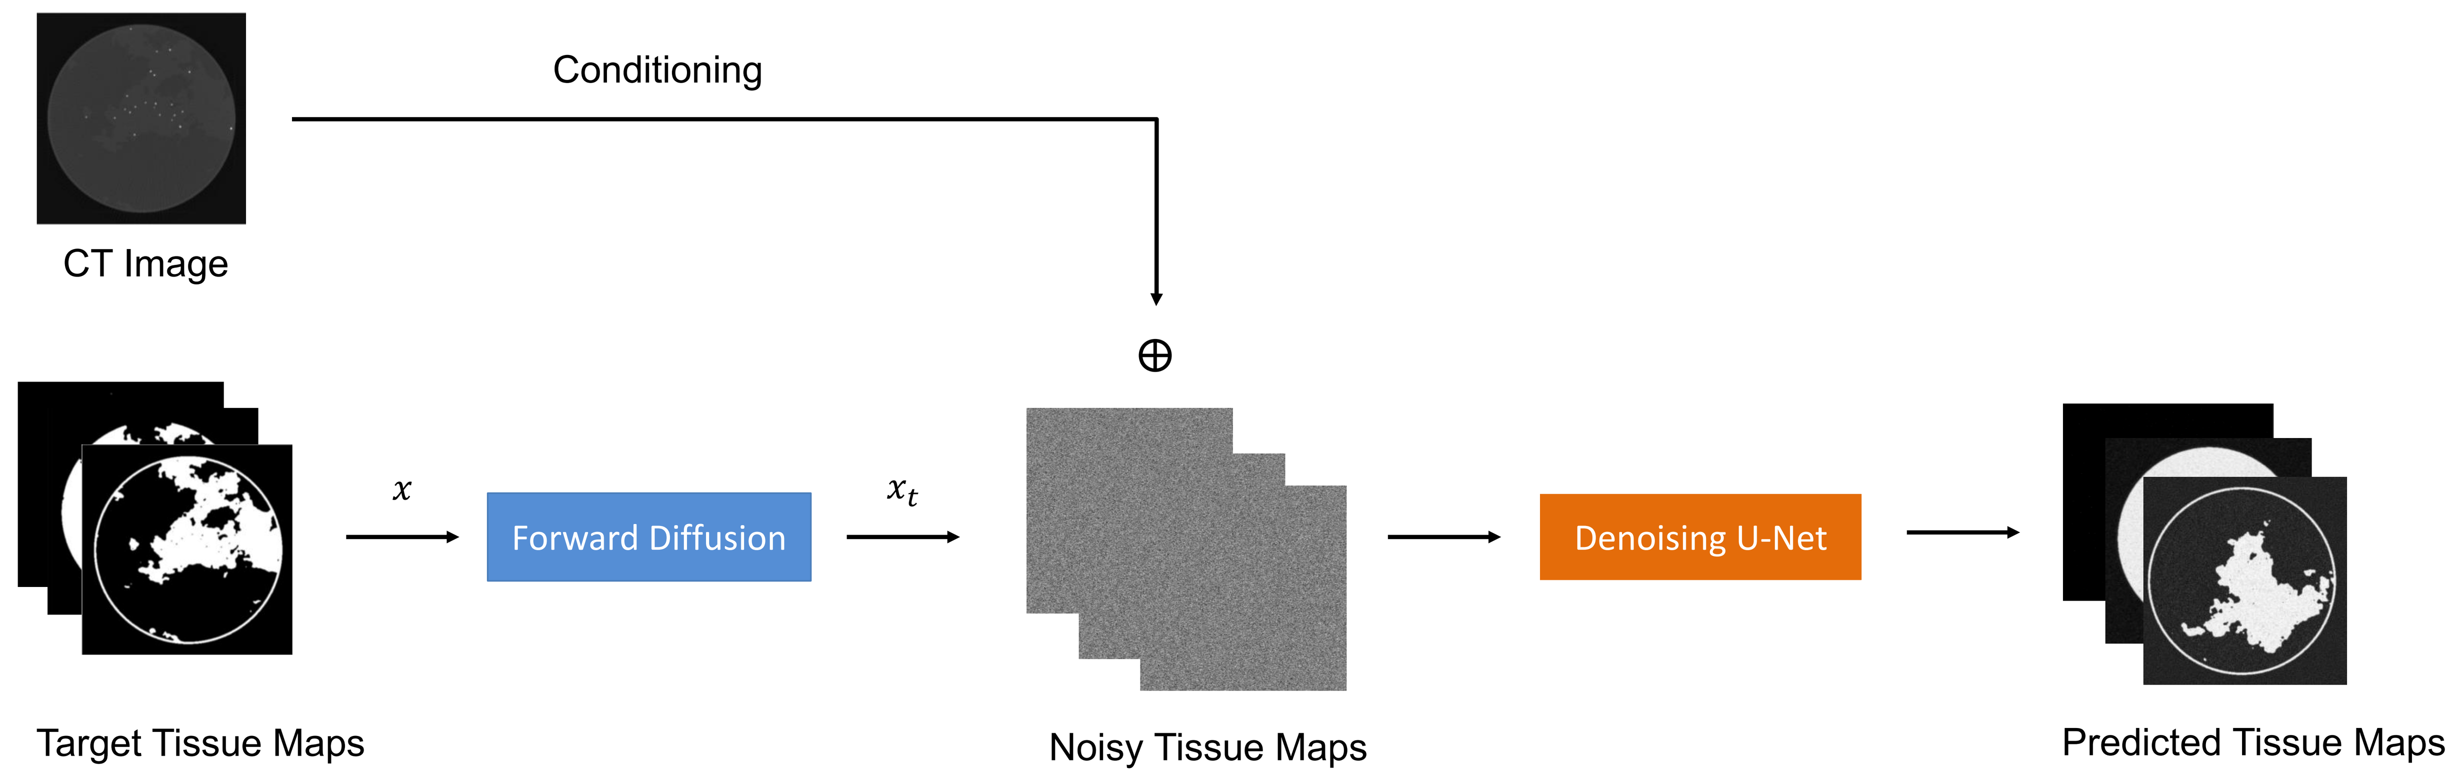

- Use U-Net backbone based on "Improved Denoising Diffusion Probabilistic Models" [5]
- Input: 4-channel (concatenation of noisy sample tissue maps and conditioning CT image).
- Output: 3-channel predicted tissue maps.
- Linear noise schedule (t=1000 steps)
- L2 loss with 5× weight on calcification
- Trained for 200K iterations
- DDIM with 100 denoising steps used for all inference.
- Applied post-processing thresholding to reduce residual noise in images

Results

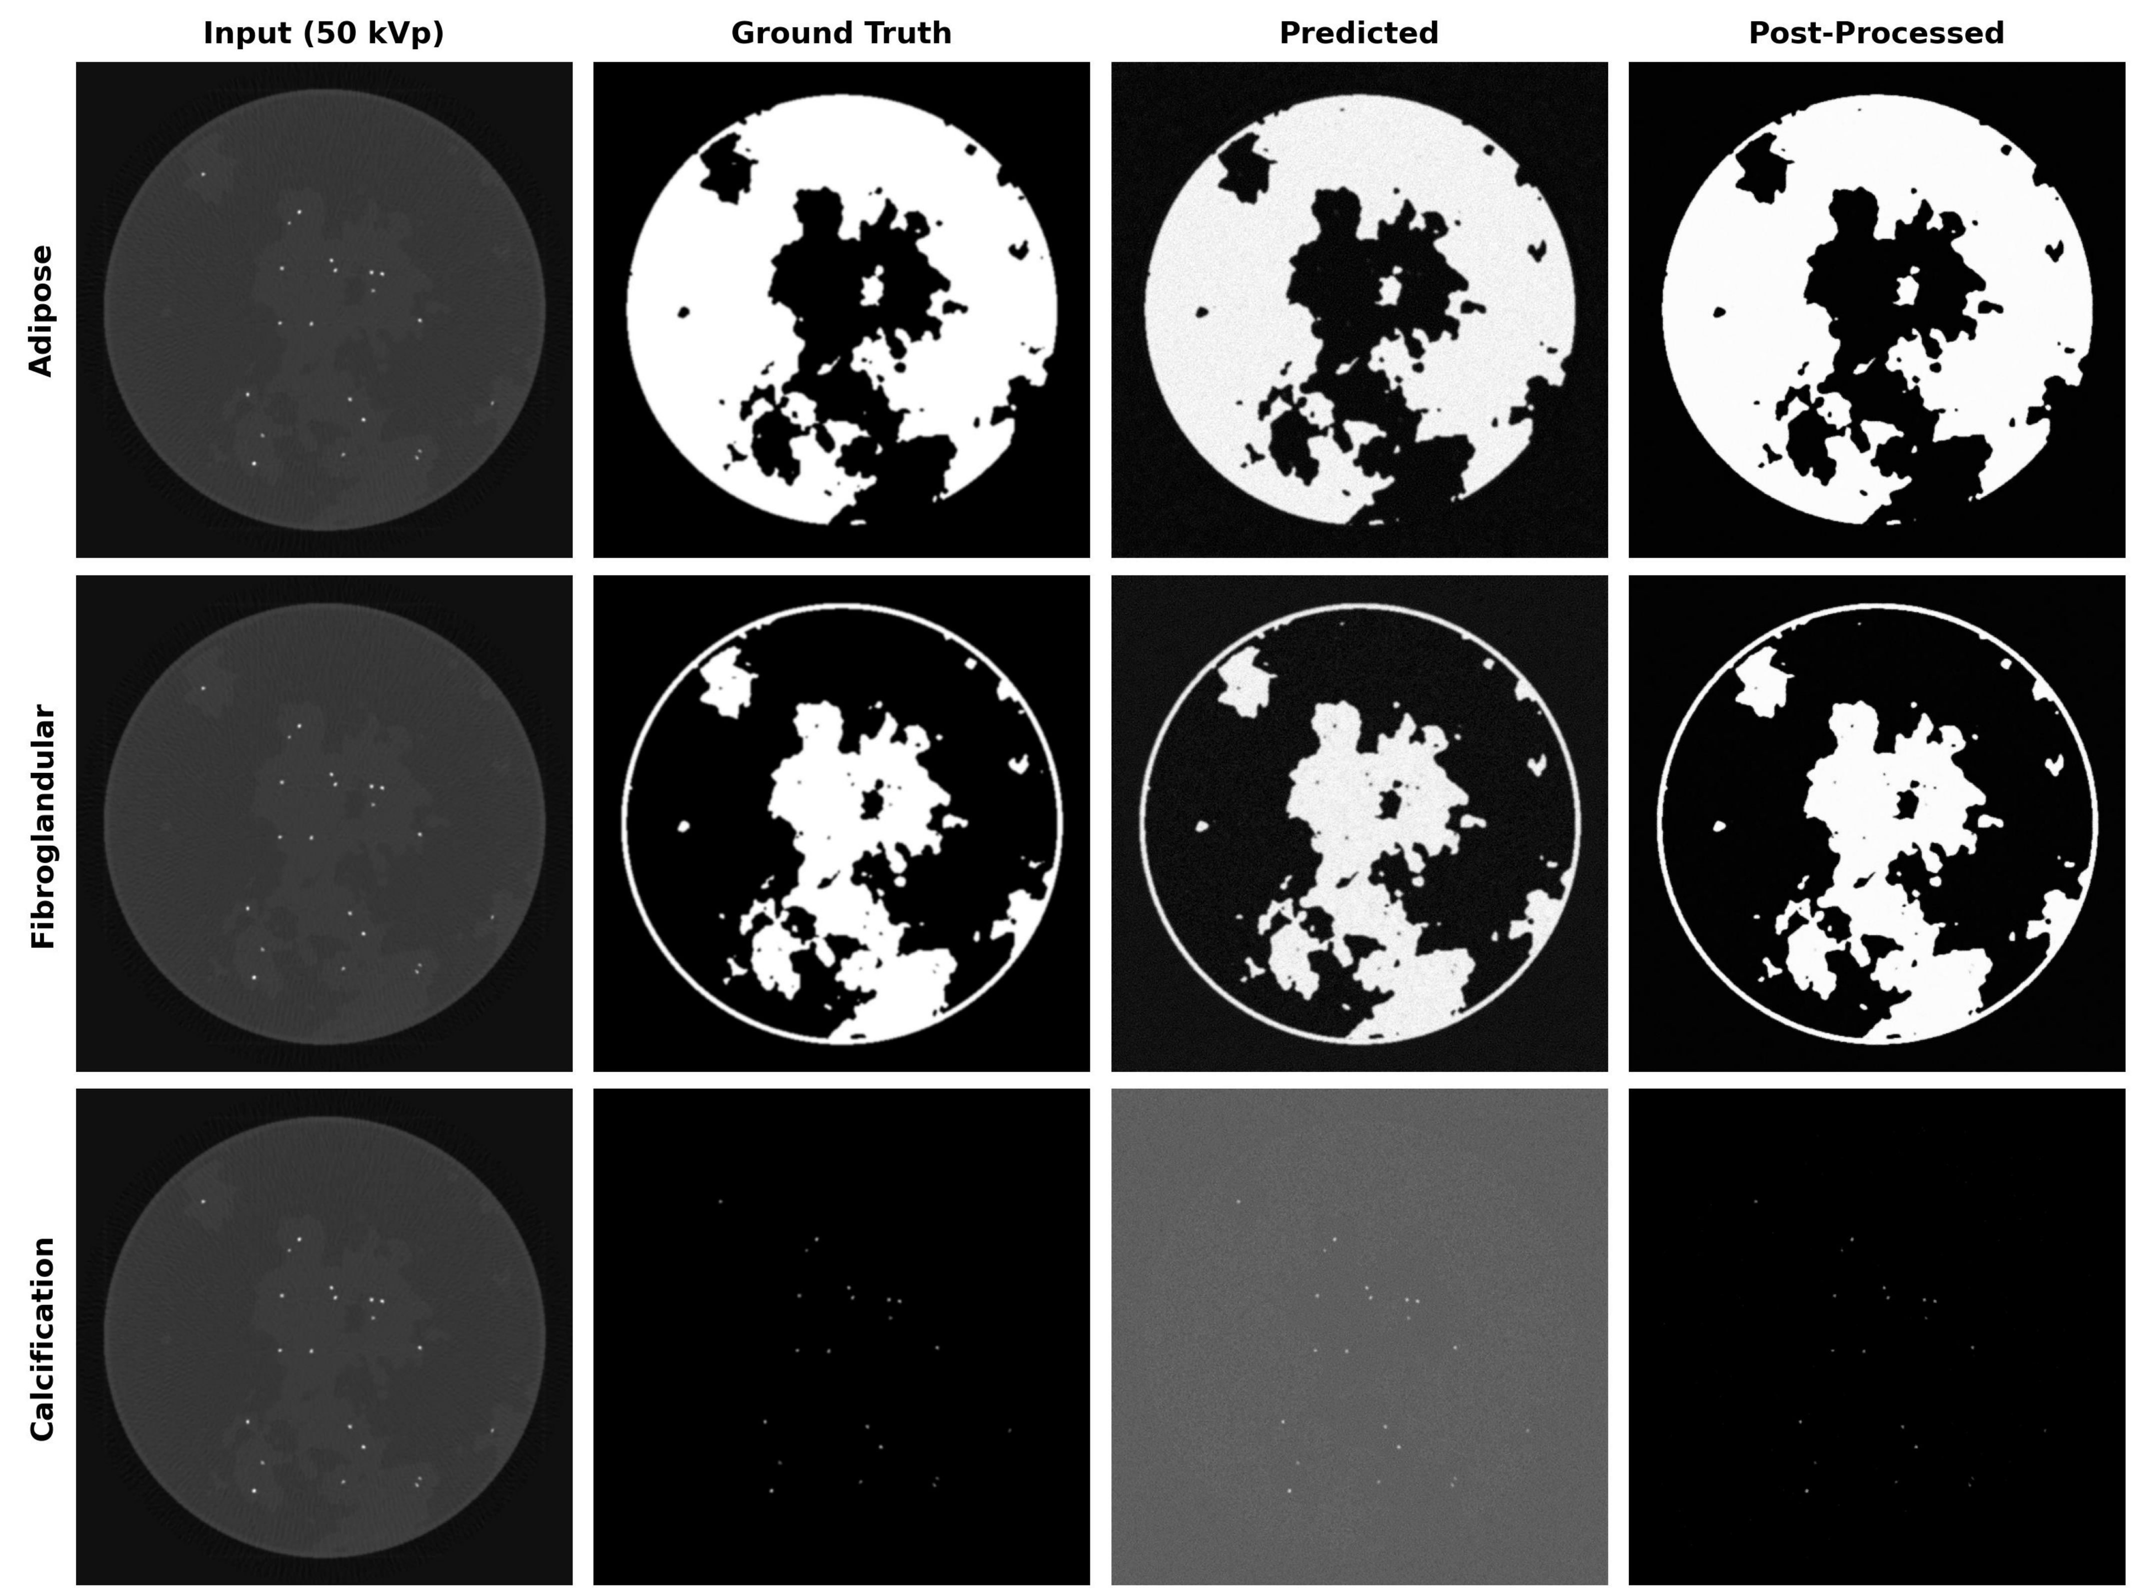

Summary Metrics (50 test samples)

Metric	Predicted Image	Post-processed
Average RMSE	0.223	0.042
Average SSIM	0.138	0.867
Average PSNR	15.5	30.1

Per-Tissue Breakdown

Tissue	RMSE (pred)	RMSE (PP)	SSIM (pred)	SSIM (PP)	PSNR (pred)	PSNR (PP)
Adipose	0.093	0.053	0.229	0.907	20.9	25.6
Fibroglandular	0.131	0.064	0.183	0.725	18.0	24.0
Calcification	0.446	0.010	0.001	0.971	7.4	40.6

Conclusion

Conditional diffusion models show promise in estimating the tissue decomposition maps from single-energy CT images for breast tissue and demonstrate the potential to approximate spectral CT information without dual-energy hardware, though residual noise and data availability remain a challenge.

Future work: Optimize model performance to produce accurate tissue maps without the need for further post-processing by fine-tuning pre-trained models to address data limitations of paired CT data.

References

[1] Sidky & Pan, "AAPM deep-learning spectral CT Grand Challenge", Med Phys, 2024
[2] Lyu et al., "Estimating dual-energy CT imaging from single-energy CT data with material decomposition convolutional neural network", Medical Image Analysis, 2021
[3] Ho et al., "Denoising Diffusion Probabilistic Models", NeurIPS 2020
[4] Gao et al., "CT-based synthetic contrast-enhanced dual-energy CT generation using conditional denoising diffusion probabilistic model", Phys Med Biol, 2024
[5] Nichol & Dhariwal, "Improved Denoising Diffusion Probabilistic Models", ICML 2021