Imagining Infinity: Endless CT Datasets through Conditional Diffusion Models

UNIVERSITY of TARTU
Institute of Computer Science
Ekaterina Sedykh, supervisor: Dmytro Fishman, PhD
BCV Lab

IDEA

Problem:
- Limited availability of high-quality medical imaging datasets.
- Logistical, ethical, and privacy concerns in acquiring real medical images.
- Class imbalance in existing datasets.
- Lack of specific cases needed (negative/specific positive cases).

Proposition:
- Utilization of conditional diffusion models to generate synthetic CT images.
- Application of inpainting technique.

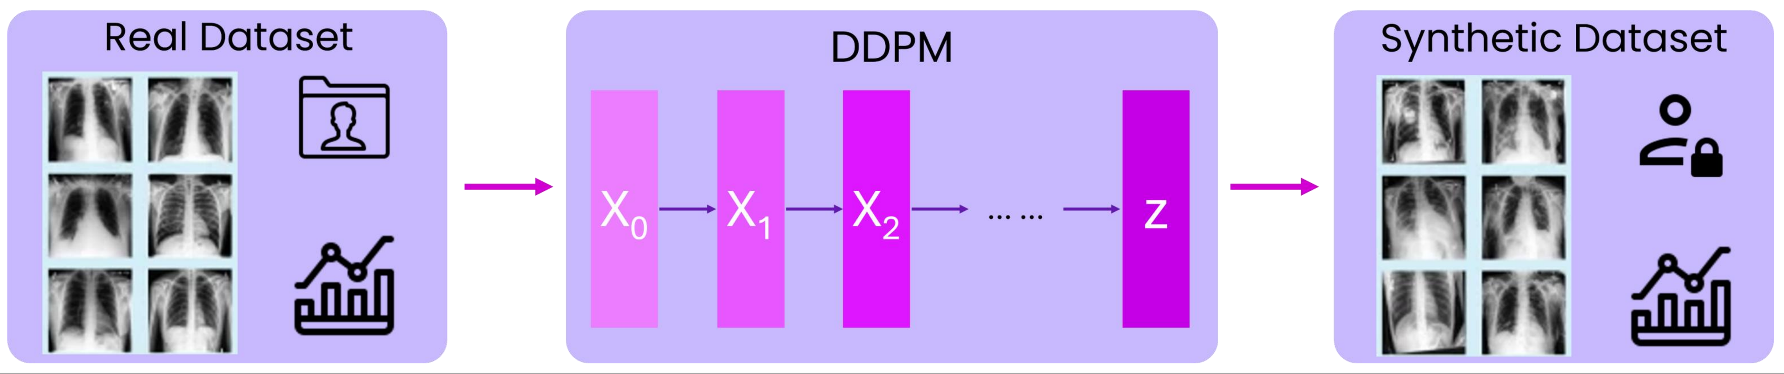

Figure 1. Generating Synthetic Medical Imaging Data Using Diffusion Models

Figure 3. 3D vs 2D data

3D CT scans are harder to work with, more complex model is required.

Solution: extract 2D slices

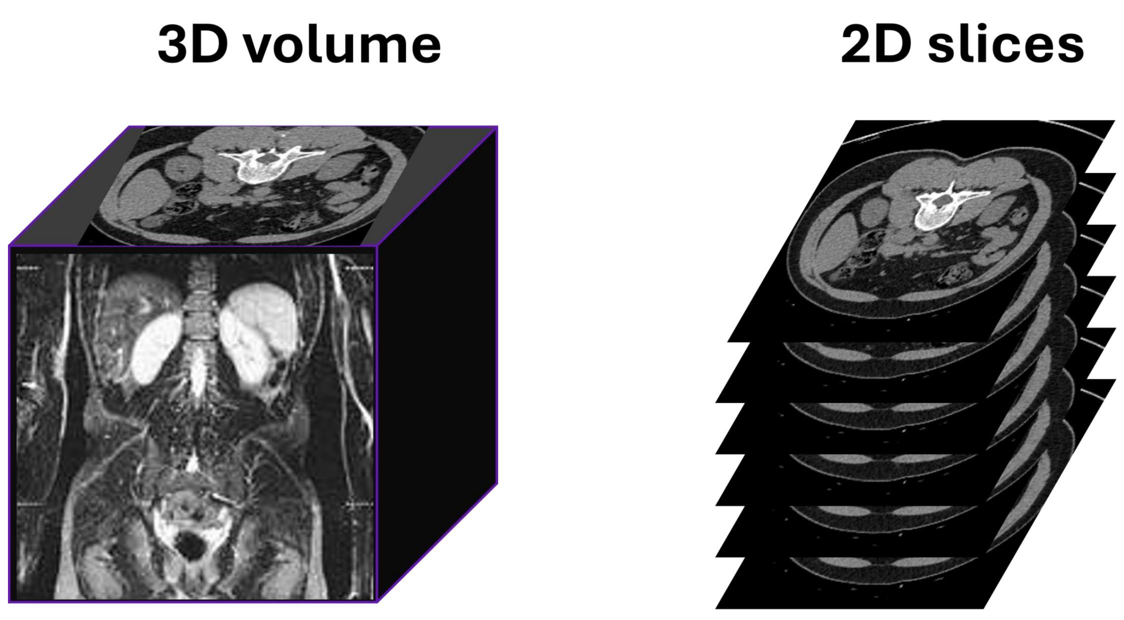

DIFFERENT OPTIONS

Inside organ boundary:
- Take the annotation of organ of interest (tried kidneys, lungs, liver, tumors)
- Can take multiple organs as well as one mask, no need for multiple channels

Outside organ boundary:
- Take the bounding box of organs of interest + margin
- Can take any other shape
- In testing no borders were notices

Generate completely new tissue:
- Trained on tumors dataset, like regular inside organ boundary task
- Create test samples by combining one CT scan with another tumor annotations -> introduce mask into a new place -> new tumor

Organ	↓FID (Pix2Pix)	↓FID (Diffusion)	↓FID (Neighbouring slice)
Kidneys	11.8897	0.4678	0.0689
Lungs	28.1489	2.4213	0.0465
Liver	19.7494	0.8220	0.0653
Tumor	15.5298	1.0933	1.2491

Table 1. FID scores for different organs generated by the diffusion model, GAN model (Pix2Pix), and benchmarking against the next slice. Lower FID scores suggest higher quality synthetic images.

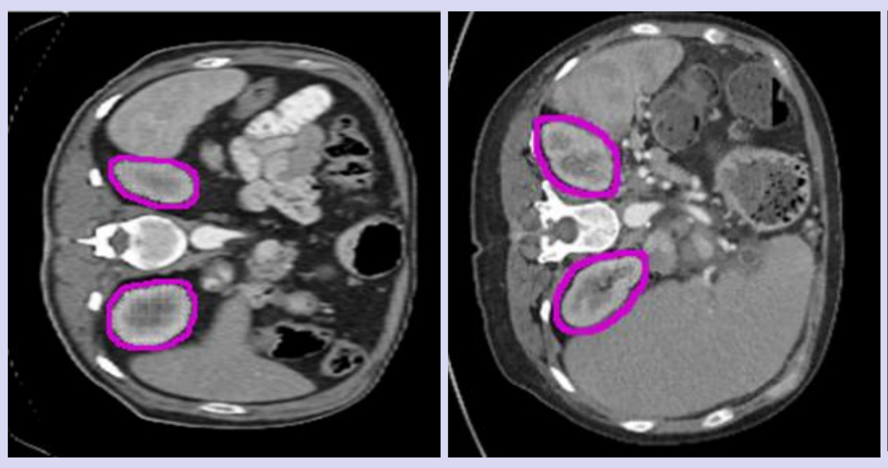

Figure 5. Example of generated kidneys by GAN model (on the left) – they are blurry and don't have kidney details, and our diffusion model (on the right)

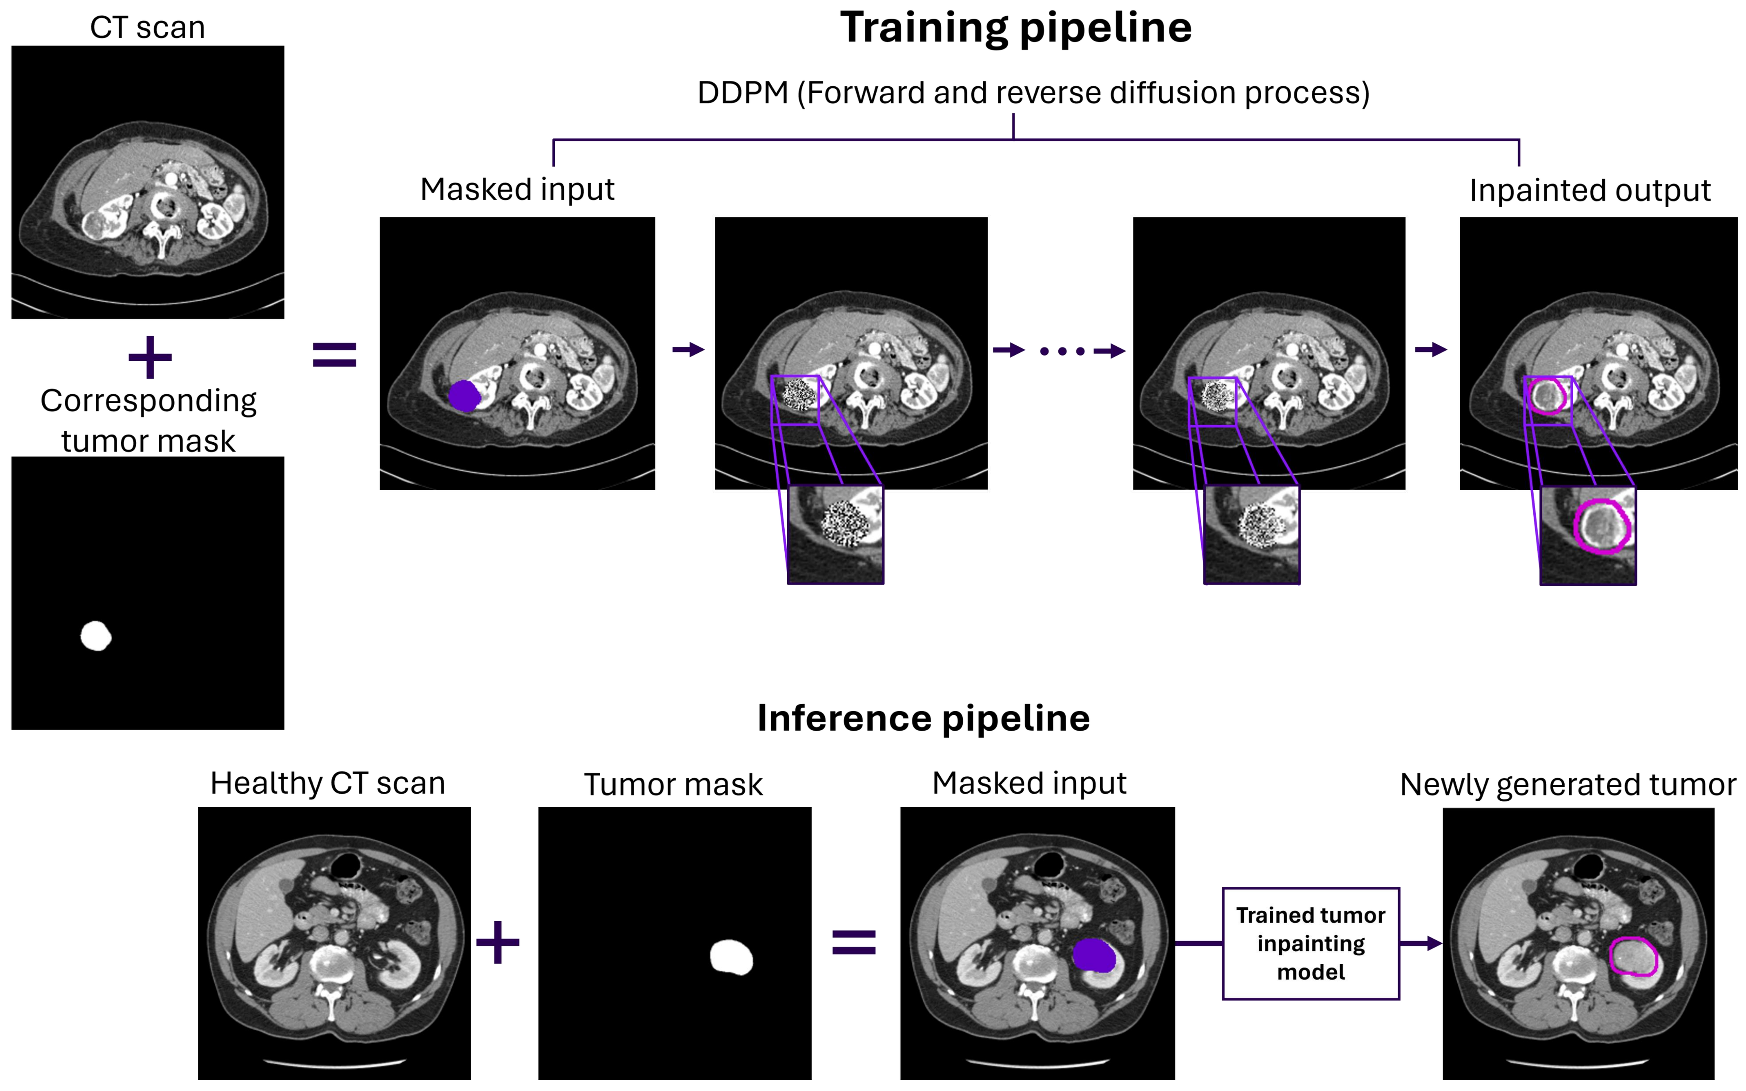

Figure 2. General training and inference pipelines for generating synthetic tumors.

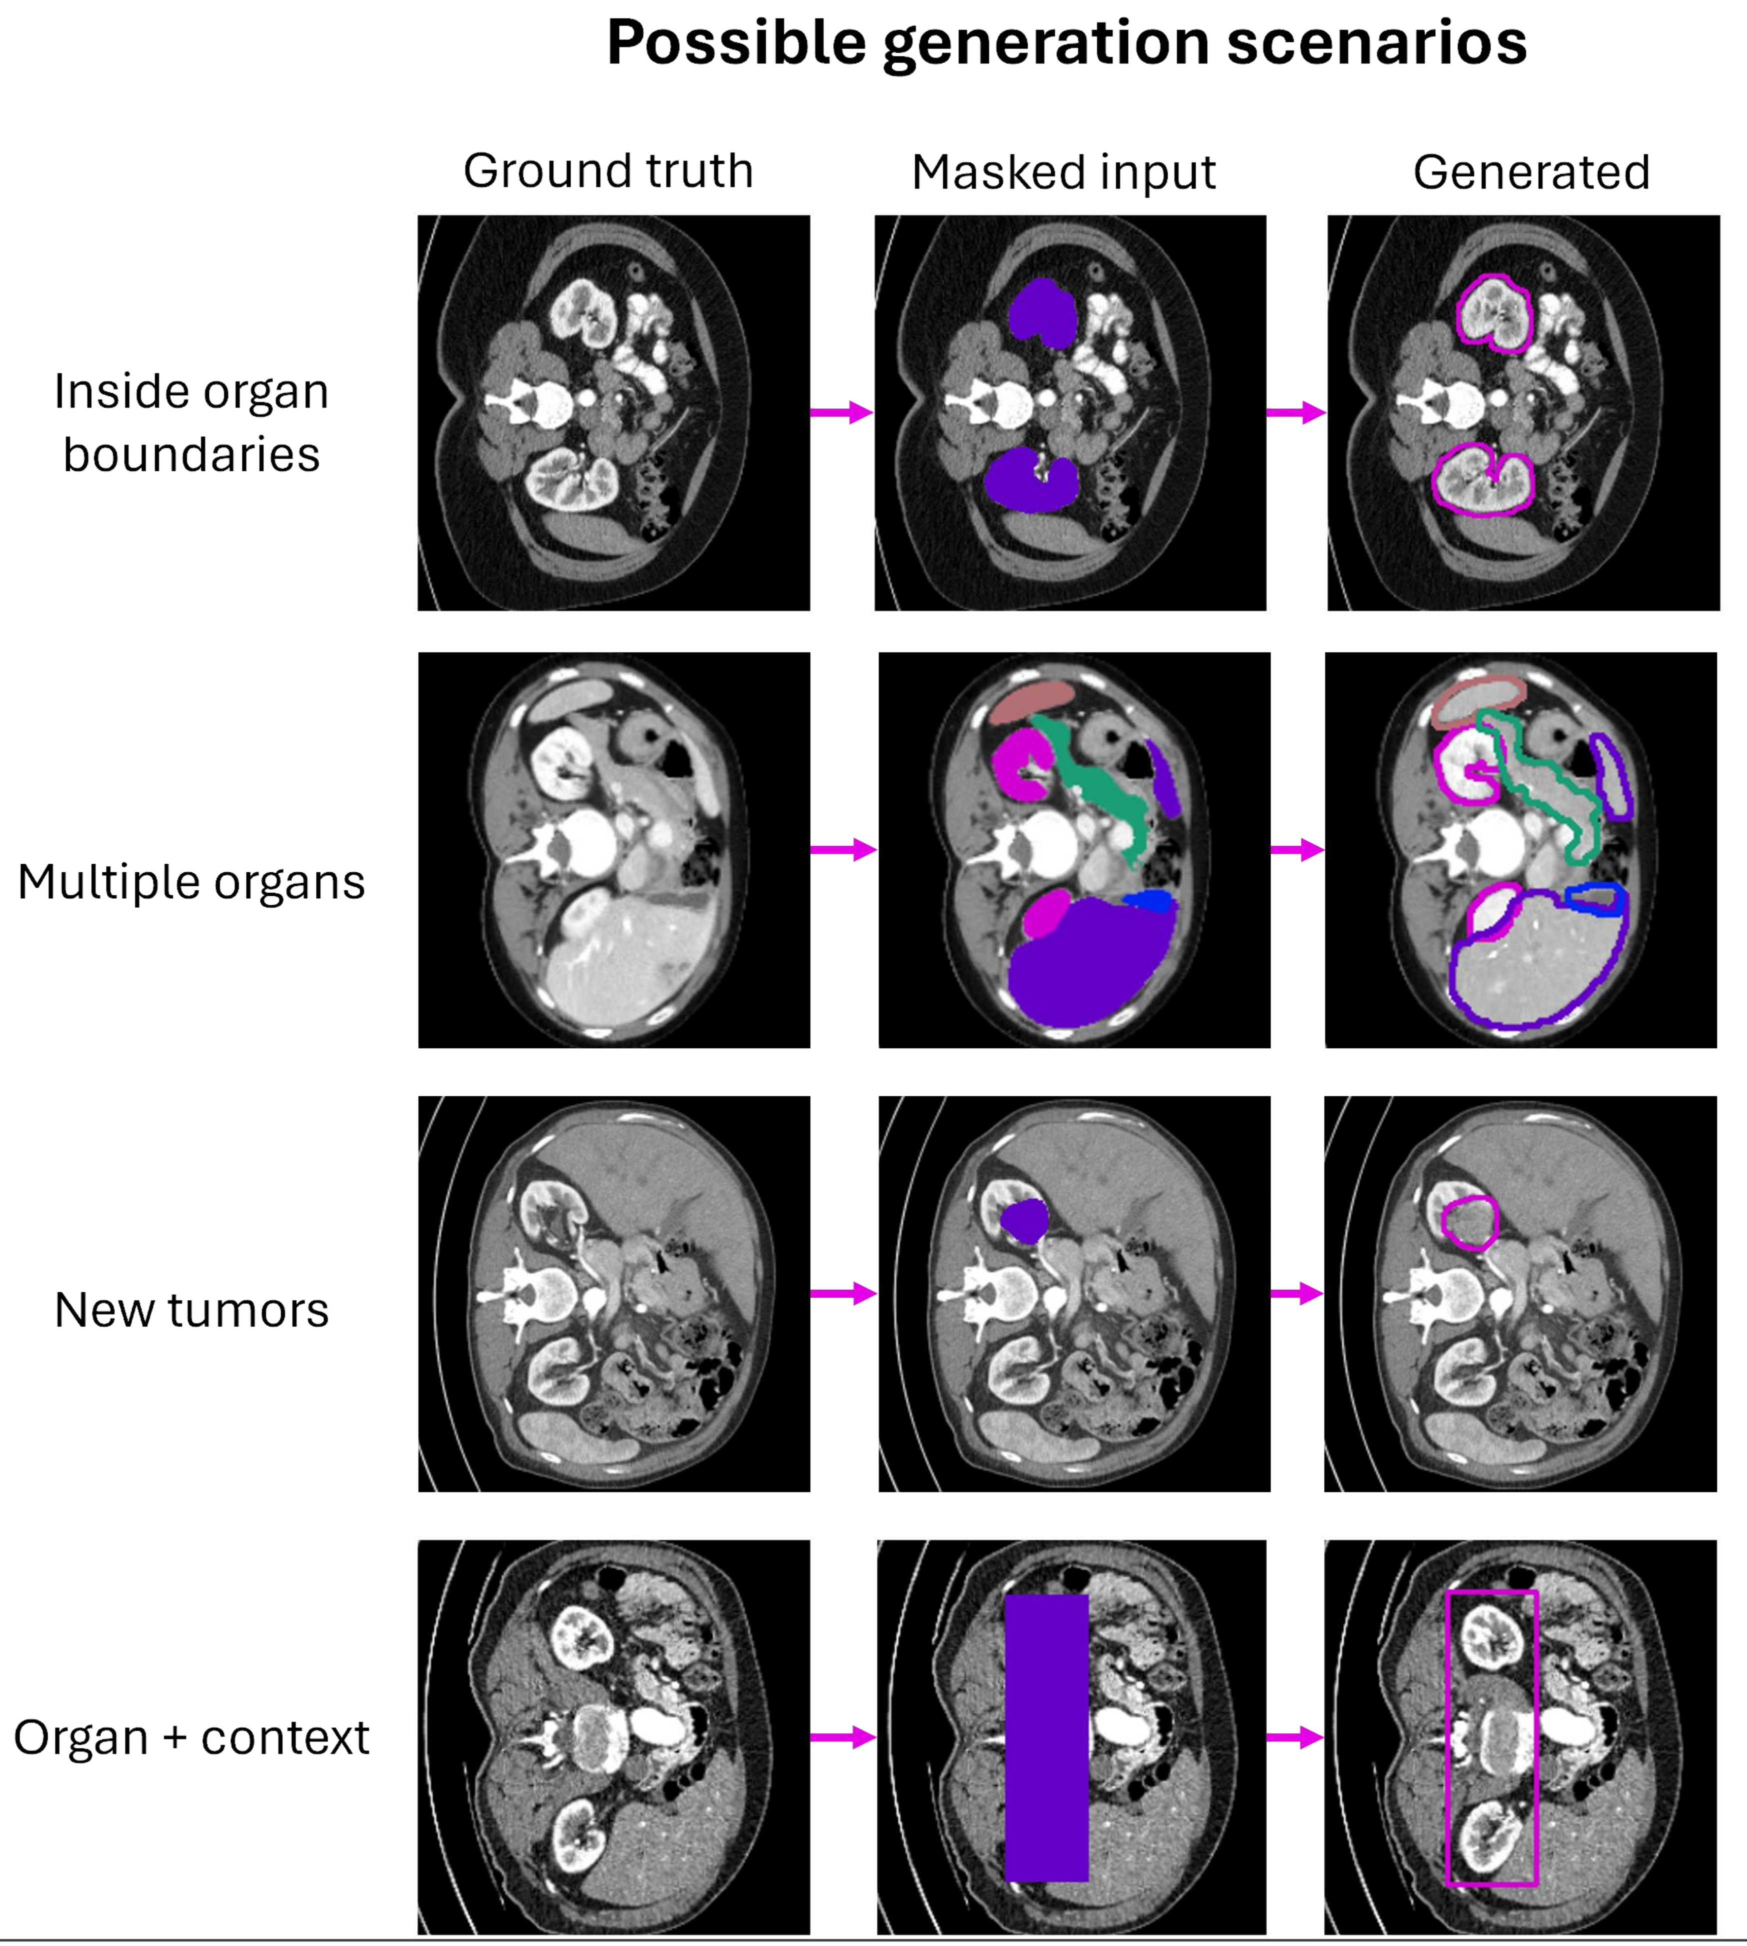

Figure 4. Model's generation capabilities: generating structures inside organ boundaries, generating multiple organs from one binary mask, generating new tumors, and generating organs with context around it. Each case includes the ground truth, masked input, and the generated output.

CONCLUSIONS AND FUTURE WORK

What we did:
- Diffusion models can generate high-quality synthetic CT scans
- Similar data generation and a new tissue (i.e. tumor) introduction
- Diffusion models consistently outperformed GANs

What's next:
- More tests on downstream tasks.
- Developing methods for 3D image generation
- Extending to other medical imaging modalities like MRI and PET