Learning CT Segmentation from Label Masks Only

Artyom Tsanda[1,2], Hannes Nickisch[3], Tobias Wissel[3], Tobias Klinder[3], Tobias Knopp[1,2,4], Michael Grass[3]

[1] Hamburg University of Technology, Institute for Biomedical Imaging, Hamburg, Germany
[2] University Medical Center Hamburg-Eppendorf, Section for Biomedical Imaging, Hamburg, Germany
[3] Philips Innovative Technologies, Hamburg, Germany
[4] Fraunhofer Research Institution for Individualized and Cell-Based Medical Engineering IMTE, Lübeck, Germany

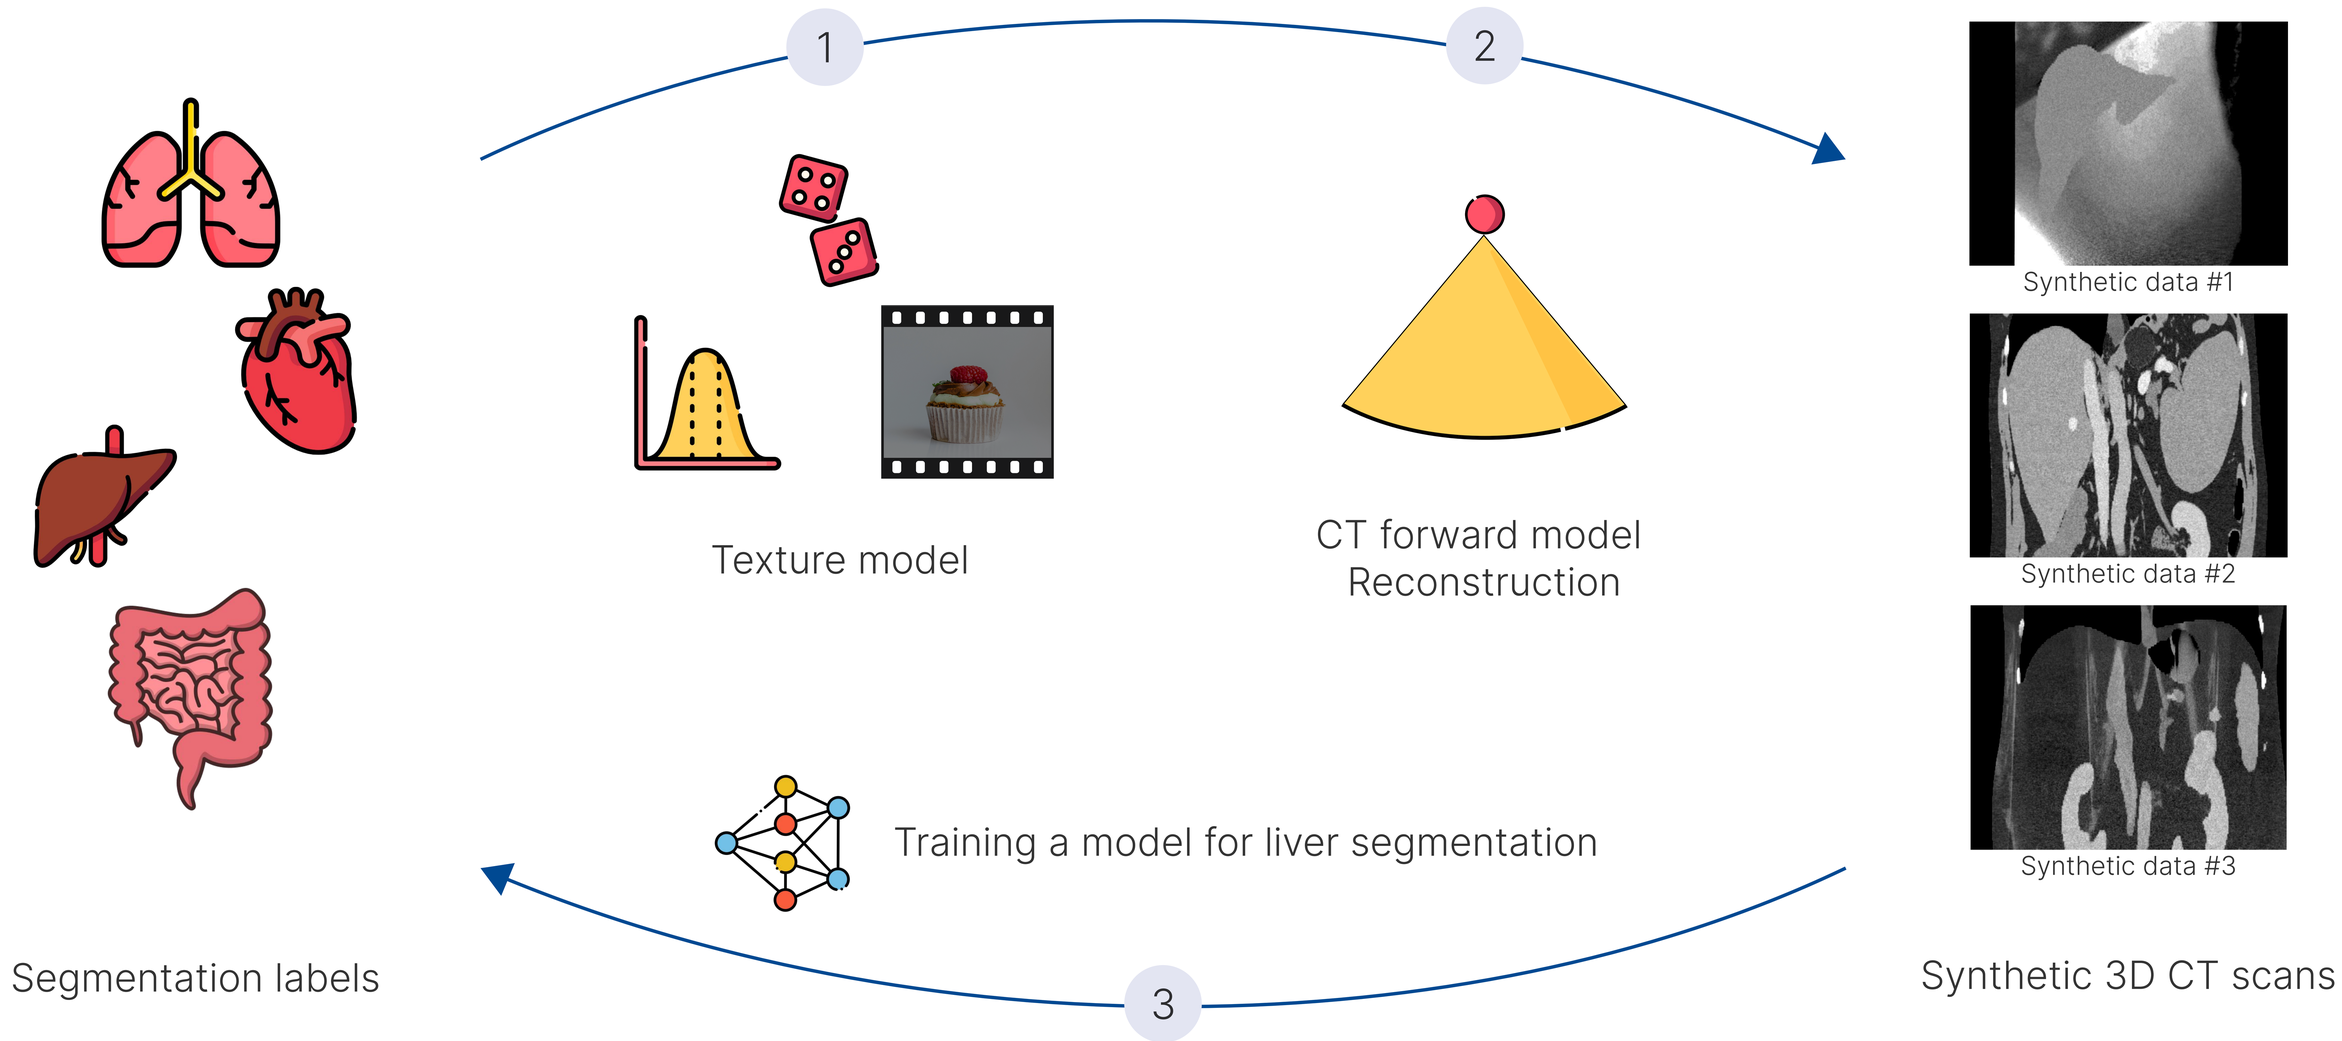

Figure 1. Proposed method for synthetic CT generation from segmentation labels.

Background

We consider the problem of training CT segmentation models with the scarcity of CT images, *i.e.* with segmentation masks only. To address the issue, one can employ generative approaches to transfer data from other modalities[1] or generate CT images directly from segmentation masks. However, these require a separate training step. We employ the approach based on heuristics to fill in the texture[2] and physics to simulate the acquisition.

Method

Data generation & training

We explore three strategies to generate synthetic CT scans from 64 segmentation masks from the TotalSegmentator dataset[3]:

- **Synthetic data #1**: the liver mask is filled with a constant value sampled from the Gaussian distribution and placed into a cylinder filled with histogram-matched natural videos
- **Synthetic data #2**: labels are filled with constant values sampled from Gaussian distributions
- **Synthetic data #3**: labels from the TotalSegmentator dataset are filled with histogram-matched videos

All scans are then forward-projected onto helical geometry to sample noise. Finally, the noisy projections are reconstructed. We employ the nnU-net framework[4] to train liver segmentation models on the generated data.

Test

We use 44 full-body CT scans from the TotalSegmentator dataset for testing.

Results

Table 1 shows averaged Dice scores for models trained on the considered synthetic datasets and a model trained on real CT scans from the training set. As the models demonstrate comparable performance, the ones trained on synthetic data successfully generalize to real CT data.

Table 1. Dice scores on the test set calculated from models trained on synthetic and real data

Training Dataset	Dice Score
Synthetic data #1	0.951
Synthetic data #2	0.971
Synthetic data #3	0.967
Clinical data	0.947

Conclusion

- We proposed several strategies to generate synthetic CT segmentation datasets that rely on segmentation labels and prior knowledge about CT images.
- The segmentation models trained on synthetic data generalize to real CT scans.
- The results indicate a promising approach to bootstrap CT segmentation data as it does not require actual images and employs an imperative strategy without the need to train proxy-models, such as GANs.
- Future work may focus on joint multi-modal data generation.

1. Zhang, Zizhao, et al. "Translating and segmenting multimodal medical volumes with cycle-and shape-consistency generative adversarial network."
2. Hoopes, Andrew, et al. "SynthStrip: skull-stripping for any brain image."
3. Wasserthal, Jakob, et al. "Totalsegmentator: Robust segmentation of 104 anatomic structures in CT images."
4. Isensee, Fabian, et al. "nnU-Net: a self-configuring method for deep learning-based biomedical image segmentation."

Corresponding Author:
Artyom Tsanda
artyom.tsanda@tuhh.de
linkedin.com/in/atsanda