Data Driven Geometrical Models of Anatomical Structures for Simulations

A PROJECT REPORT
SUBMITTED IN PARTIAL FULFILMENT OF THE
REQUIREMENTS FOR THE DEGREE OF
Master of Technology
IN
Faculty of Engineering

BY

Aditya Kumar

Computer Science and Automation
Indian Institute of Science
Bangalore – 560 012 (INDIA)

July, 2022

Declaration of Originality

I, **Aditya Kumar**, with SR No. **04-04-00-10-42-20-1-18235** hereby declare that the material presented in the thesis titled

Data Driven Geometrical Models of Anatomical Structures for Simulations

represents original work carried out by me in the **Department of Computer Science and Automation** at **Indian Institute of Science** during the years **2021-2022**.
With my signature, I certify that:

- I have not manipulated any of the data or results.
- I have not committed any plagiarism of intellectual property. I have clearly indicated and referenced the contributions of others.
- I have explicitly acknowledged all collaborative research and discussions.
- I have understood that any false claim will result in severe disciplinary action.
- I have understood that the work may be screened for any form of academic misconduct.

Date: Student Signature

In my capacity as supervisor of the above-mentioned work, I certify that the above statements are true to the best of my knowledge, and I have carried out due diligence to ensure the originality of the report.

Advisor Name: **Vijay Natarajan** Advisor Signature

DEDICATED TO

My parents

Acknowledgements

First and foremost, I thank my advisor, professor Vijay Natarajan for all the guidance and support he provided throughout the course of the project. Due to the pandemic and the consequent absence of face-to-face interaction, regular online meetings with him helped me stay motivated and keep working in the right direction.

I am also very grateful for the help and support provided by Nithin Shivashankar of Mimyk, who was a like a second advisor for me. He provided a lot of valuable insight and critique throughout my work, as well as suggestions for directions I should work on.

I appreciate the inputs from Dr. Kirthi Koushik from the HCG Ramaiah Cancer Center. He helped me a lot in understanding the Brachytherapy procedure at Ramaiah, the technology they use, and the limitations they face. I also thank miss Revati for painstakingly going through the patient database at Ramaiah and extracting all the data I required for this project.

Finally, I would like to thank my friends at the Indian Institute of Science and my family for being a constant source of motivation and support.

Abstract

Brachytherapy is a cancer treatment methodology based on highly localised delivery of radiation to tumor regions. It ensures minimal damage to surrounding tissue and organs and provides higher chance of long-term survival as well as lower recurrence rates in comparison to other radiation based treatments. This project explores methods for creating geometrical models of anatomical structures for assisting in Brachytherapy planning as well as for creating simulations useful for training new practitioners.

Current work in the project is focused on segmenting out high-risk organs from a CT scan of a patient automatically or semi-automatically, which is an important step in planning the treatment. The goal is to create a fully automatic or interactive tool for contouring said high-risk organs which can be integrated into the treatment workflow.

Contents

List of Figures

Chapter 1

Introduction

Modern cancer treatment methods make extensive use of radiation therapy. Brachytherapy is prominent among such methods for delivering highly targeted radiation to tumor regions while avoiding damage to surrounding organs and tissues. The high precision allows for the ability to expose tumor regions to high radiation dosages in a short time, making Brachytherapy a relatively quicker alternative to other methods like External Beam Radiation Therapy (EBRT). Brachytherapy can also be used as a complementary treatment to boost radiation dosage from other treatment methods. Overall, Brachytherapy offers a relatively safer and surer method for treatment of cancers in a multitude of tumor locations (head and neck, cervix, prostate, breast, skin etc.)

1.1 Treatment Methodology

As stated earlier, Brachytherapy is based on highly localised delivery of radiation to tumor regions. Delivery is made via permanently placed seeds containing radioactive material, or via afterloaders that can deliver high-dose radiation in an unsupervised manner after proper setup. This ensures minimal damage to surrounding tissue and organs and provides a higher chance of long-term survival and lower recurrence rates. Brachytherapy can be classified as follows:

- Based on duration of radiation.
  - **Permanent**: Very low dose seeds left permanently in body.
  - **Temporary**: Higher dose sees positioned by afterloaders deliver radiation for a period of time, and can be removed after treatment.
- Based on position of radionuclides.

  - **Interstitial Brachytherapy**: Radiation sources are placed inside the tumor.
  - **Contact Brachytherapy**: Radiation sources are placed close to tumors. Includes intracavitary and surface Brachytherapy.

- Based on dose-rate.
  - **Low Dose Radiation (LDR)**: In patient procedure. Involves a long continuous delivery of low dose radiation.
  - **High Dose Radiation (HDR)**: Usually out-patient procedure. High dose is delivered in a short time.
  - **Pulsed Dose Radiation (PDR)**: Slightly lower dosage than HDR. Applied multiple times (pulses of $\sim$10 minutes).

1.2 Brachytherapy in India

Brachytherapy has a long history in India, its first documented usage being in 1920 [14]. Currently, HDR Brachytherapy is the norm, with around 250 HDR Brachytherapy centers across India. The major use of Brachytherapy in India is for treatment of cervical cancer, of which India gets a huge number of cases per year. GLOBOCAN's 2018 estimate shows nearly 100,000 new cases for cervical cancer annually, with nearly 60,000 of them resulting in deaths [4].

Other than cervical cancer, head and neck cancers form another large portion of cancer cases in India. Among women, breast cancer is the most common after cervical cancer. Somewhat contrary to the focus on prostate cancer in literature, prostate cancer rates in India are much lower than America (9-10 in 100,000 vs 106.5 in 100,000 [14]). Hence, prostate Brachytherapy is not practiced much.

As for the planning methodologies in use, MRI guided Brachytherapy is generally infeasible due to high costs (MRI guidance is useful correct position of radiation sources). Instead, using CT and ultrasound scans is more prevalent. Studies show that overall survival and toxicity rates observed using CT scans and ultrasounds is close to that offered by MRI guided Brachytherapy.

1.3 Challenges and Opportunities

The major challenge in India is the lack of infrastructure to handle all cases. Further, the current reimbursement models for Brachytherapy end up making other treatment methods more lucrative for practitioners. Finally, there is a severe lack of skilled professionals for procedures like needle insertion and applicator placement. These challenges present us with the following opportunities:

- It is possible to speed up the treatment process by speeding up the time consuming planning phase using computational techniques.
- To train new practitioners, there is a need to create simulation systems that help trainees to practice on patient specific simulations instead of practicing on actual cases, which is extremely risky.

Chapter 2

Problem Description

Owing to the prominence of cervical cancer cases in India, and the need for segmenting out organs from 3D volume data for both treatment planning and simulation, work so far has focused on segmenting out organs at risk for cervical cancer Brachytherapy, namely the bladder, rectum and sigmoid.

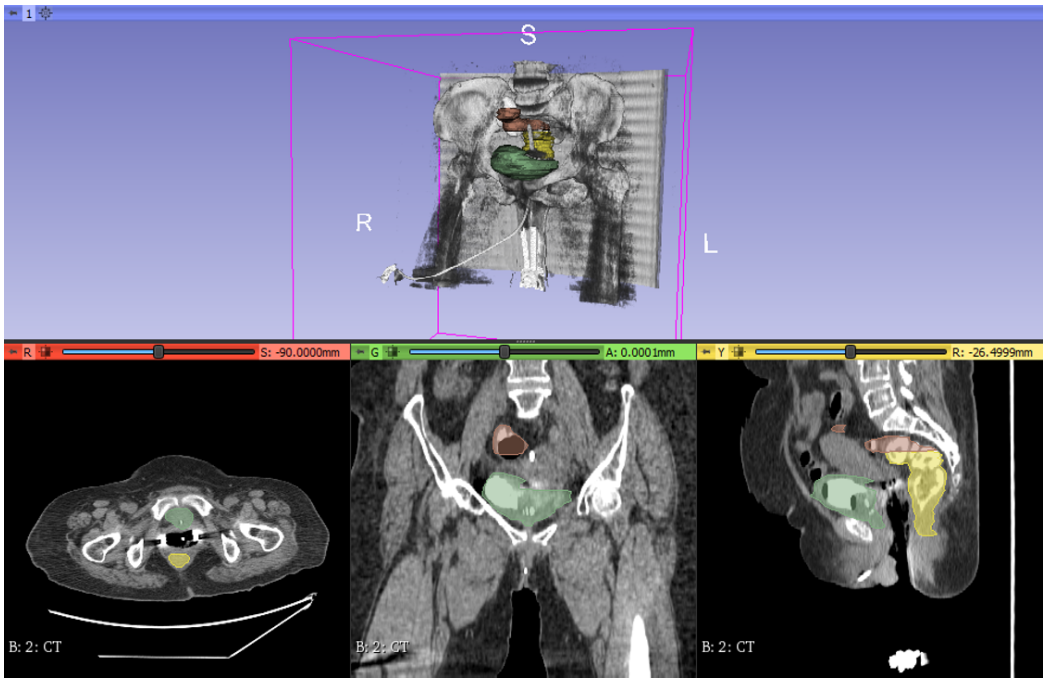

Figure 2.1: Example segmentation of CT-scan volume. Upper view shows a volume-rendering of the scan, along with the segmented organs. Lower view shows segmentations on selected 2D slices along the axial, coronal and sagittal planes. The segmented organs are bladder (green), rectum (yellow) and sigmoid (brown). *Visualization generated with 3D Slicer*

2.1 Formal Statement

Given an input tensor X of size $p \times q \times r$ with scalar values at each voxel (representing the CT volume), the aim is to output another tensor T of size $p \times q \times r$, each of whose elements is the correct label identifying the corresponding element in X to be part of one of a predefined set of regions. For the cervical cancer scenario we are considering, there are 4 possible labels: "bladder", "rectum", "sigmoid" and "none of these".

2.2 Contributions

- Data collection – 132 CT scans collected from HCG Ramaiah Cancer Center, anonymized and converted to suitable formats for easy analysis and model training.
- Baseline automatic segmentation model with UNet3D.
- Investigation of various losses for automatic segmentation.
- Training and testing a 2D interactive segmentation method – InterCNN
- Design, implementation and evaluation of a 3D extension of InterCNN.

Chapter 3

Related Work

Medical image segmentation is a challenging problem due to unclear boundaries and variations in imaging methodologies. There has been a lot of work on both automatic and semi-automatic/interactive methods.

3.1 Automatic Segmentation

The original U-Net [23] by Ronneberger *et al* was a CNN-based deep neural network for segmenting 2D images. It used a U-shaped encoder-decoder architecture along with skip connections. It was later extended in work by Çiçek *et al* to a 3D version called 3D U-Net [7]. A very similar 3D extension architecture in the form of V-Net [19] was proposed around the same time by Milletari *et al*. Later automatic segmentation methods include HighRes3DNet [15] - which is a CNN-based neural network focused on efficient usage of CNN building blocks like residual connections and dilated convolutions to make a more compact network, MultiPlanarUNet [22] - which aims to be a general purpose, efficient and easy to use method by learning to segmenting slices of volumes along random axes and combining them to get better segmentation, and nnFormer - which experiments with using the popular Transformer architecture for segmentation tasks.

3.2 Interactive Segmentation

Besides automatic segmentation, there has been a lot of work attempts towards utilizing interactive inputs from a human user in order to achieve more accurate segmentation results. Early methods include the the work by Boykov *et al* [3] where they utilize graph-cut algorithms to find optimal segmentations for n-dimensional images. Another early method is GeoS [9] by Criminisi *et al* where they utilize Geodesic distances to pose and solve segmentation as an en-

ergy optimization problem. More modern, neural network based approaches include InterCNN [5] by Bredell *et al* - which trains a CNN to update a given segmentation on the basis of user interactions, DeepIGeoS by Wang *et al* - which incorporates deep learning with the previous GeoS framework, work by Lin *et al* [16] - which seeks to utilize the first click on a 2D image as a guide for more accurate segmentations, and work by Liao *et al*, which proposes to use multi-agent reinforcement learning to train a model to segment interactively.

Chapter 4

Data Collection and Preprocessing

All of the data so far has been collected from the HCG Ramaiah Cancer Center, Bangalore. Although the persons in charge at the HCG Centre have agreed to provide access to all data available, due to software issues, and concerns about anonymity of patients, data collection has been relatively slow. At first, 50 total cases (CT scans along with contour/segmentation data) were collected, all in DICOM format. The results on the baseline model are based on this smaller dataset. The dataset has been since expanded to 132 samples, of which 100 were used for training, 12 for validation and 20 for testing. This split was decided randomly and remained fixed for the rest of the duration of the project.

4.1 Preprocessing

Due to an inconsistency in naming of patients in CT scans and corresponding contour data (DICOM RT format), the data had to be processed to fix the inconsistency. Further, all patient identifying fields were scrubbed or replaced in order to anonymize the data. A python script based on the source code from [20] was created for this purpose.

After anonymization, the data had to be converted to formats more standard in machine learning applications - namely NumPy and Nifti. This was accomplished using the tool provided in [1].

Chapter 5

Baseline Model

As a baseline, a 3D-UNet [7] was selected for the segmentation task. Due to the large volume size (around 512x512x120, varying across examples), training the network with whole volumes failed. As a result, training with patches was employed.

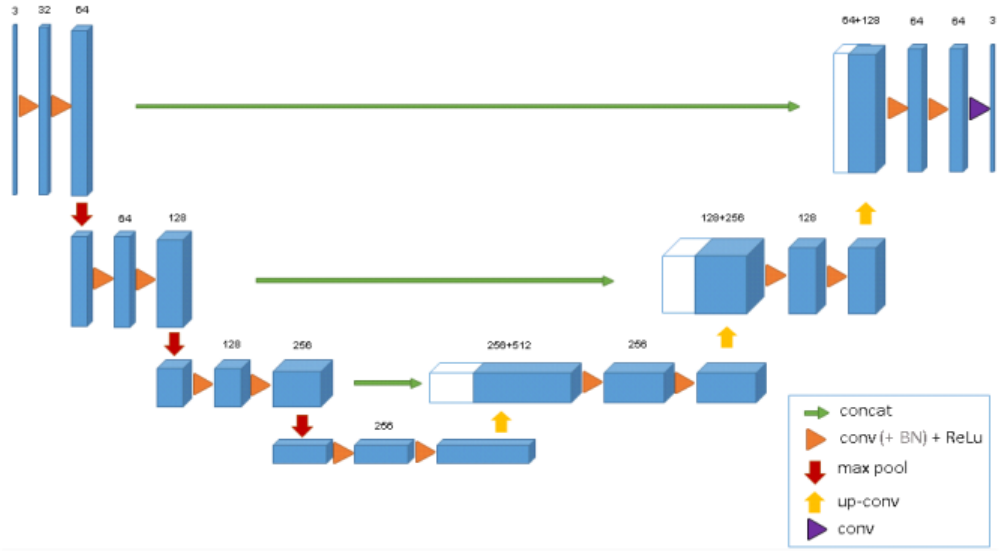

Figure 5.1: The 3D-UNet architecture. Blue boxes represent feature maps. The number of channels is denoted above each feature map. (*Source: [7]*)

The TorchIO library [21] was used to efficiently load in data, sample patches, and create data loaders for use with the PyTorch implementation of 3D-UNet network. The loss function used for training was the DICE Score (or rather, 1 - DICE Score), which is defined as

$$\text{DICE Score} = \frac{2 \times |X \cap Y|}{|X| + |Y|}$$

where X and Y represent either of target and predicted segmentations.

5.1 Choice of Loss Function

Besides the DICE score, another very similar and widely used loss for segmentation tasks is the Jaccard similarity, or the Intersection over Union (IoU) score, defined as follows.

$$\text{IoU} = \frac{|X \cap Y|}{X \cup Y}$$

There isn't a clear choice between DICE score and IoU. It can be shown that they are always positively correlated (if, on a given sample, a classifier achieves higher IoU than another classifier, then it will also necessarily have higher DICE score, and vice versa). In fact, the IoU is bounded as DICE/2 $\leq$ IoU $\leq$ DICE [2].

5.2 Training Details

The following augmentations were applied to the data before being fed into the training pipeline:

- Histogram normalization
- Z normalization
- Crop/padding the volume to size 256x256x128
- Random motion artifacts
- Random bias field (inhomogeneity introduced by variations in a magnetic field)
- Random noise
- Random flipping
- One of a random affine transformation or elastic deformation.

Due to unexpected duplicacies in the collected dataset and the unavailability of all three organs under consideration in all the samples, the data usable for training was reduced to 38 samples. 30 of these were used for training, while 8 were used for validation. Batch size for both training and validation was fixed to 2.

During training, a fixed number of patches (4) were randomly sampled from each batch and the model ran on those. Similarly for validation.

5.3 Training Results

After testing with smaller patch sizes ($32 \times 32 \times 32$ and $64 \times 64 \times 64$) the model was unable to get a satisfactory score either in training or validation data (DICE score was stuck around 0.25 after a while). This was possibly due to the patches being too small to capture patterns in a whole segment (corresponding to either the bladder, rectum or sigmoid). Increasing the patch size to 128^3 alleviated this, and the model was able to obtain a better average DICE score of 0.686 on the validation set (after 35 epochs of training). We also experimented with resampling the CT volumes to 128^3 instead of using patches, and along with the larger dataset, this improved the validation DICE score to 0.7.

Chapter 6

More Losses

In order to improve the results of the baseline model, we tried a few alternative losses.

6.1 Generalized DICE

The generalized DICE score, introduced in the work by Sudre *et al* [25] aims to weight the contribution of each class to the overall loss based on the relative volume occupied by that class. It is calculated as follows:

$$\text{Generalized DICE} = 2\frac{\sum_{l=1}^{n} w_l \times (|X_l \cap Y_l|)}{\sum_{l=1}^{n} w_l \times (|X_l| + |Y_l|)}$$

Where n is the number of classes, w_i are corresponding weight, and X_l, Y_l are respectively, correct and predicted segmentation for class l. Here, w_l were calculated as the inverse of volume occupied by each class, thus giving highest weightage to classes that occupy the least volume.

This score ended up having a critical flaw when used as a loss for segmentation training. It is possible to get a "good enough" generalized DICE score by predicting the whole volume to be the background. While training, this lead the network learning to do exactly this, and then being unable to improve further, with a generalized DICE of around 0.4. To alleviate this, we tried to "warm start" the network by training it with the usual DICE until it could do better than 0.4 generalized DICE score, and then trained it with the generalized DICE score instead of the usual DICE score. Results are shown in the following figure.

Figure 6.1: Training results with generalized DICE loss and a "warm started" 3D UNet

6.2 Weighted DICE

This loss was simply a variant of the usual DICE where the DICE scores, calculated for each class separately, were given weights and added together to form the final loss. Let n be the number of classes, $w_1, w_2, \ldots, w_n \in (0, 1)$ be weights assigned to them, and $D_1, D_2, \ldots, D_n$ be the corresponding DICE scores. The weighted DICE score is calculated as

$$\text{Weighted DICE} = w_1 D_1 + w_2 D_2 + \ldots + w_n D_n$$

The 3D UNet was trained with similar a similar training setup as before. To improve the lower DICE scores on the Rectum and Sigmoid classes, the weights were assigned as follows:

- **Background**: 0.05
- **Bladder**: 0.2
- **Rectum**: 0.3
- **Sigmmoid**: 0.45

Even with these highly skewed weights, training results ended up being the same as in the case of usual DICE loss. That is, the overall DICE score on the validation set turned out to be the same as before, and there was no improvement in the scores for Rectum and Sigmoid classes.

6.3 Topological Loss

With the baseline model, we noticed that in a few cases, none of the voxels in the volume were marked as belonging to a some classes. For example, in one output segmentation, none of the voxels were marked to be belonging to Sigmoid, even though the Sigmoid was supposed to be present in the volume. This motivated us to incorporate a loss that would make sure that the topology/shape of the output segmentations is correct (most importantly, the number of components).

We used the loss described in the work by Clough *et al* [8]. The basic idea is to calculate the persistence homology of the super level sets of each segmented object, then comparing the calculated homology to the expected Betti numbers for the object. The persistence barcode, listing the birth and death point of each topological feature, is used to determine the most persistence features, make sure they add up to the required Betti numbers, and minimize the persistence of other features.

If β_k^* denotes the expected k-dimensional Betti number, $b_{k,l}$ and $d_{k,l}$ denote the birth and death values for l^{th} most persistent k-dimensional feature from the persistence barcode, then the topological loss can be calculated as follows:

$$\mathcal{L}_k(\beta_k^*) = \sum_{l=1}^{\beta_k^*}(1 - |b_{k,l} - d_{k,l}|)^2 + \sum_{l=\beta_k^*}^{\infty}|b_{k,l} - d_{k,l}|^2$$

$$\mathcal{L}_{topo} = \sum \mathcal{L}_k(\beta_k^*)$$

This enforces exactly β_k^* bars of length equal to 1 (assuming the filtration for persistence calculation is performed on a probability distribution predicted by a model). During training, this loss was used in combination with the usual DICE loss, with a weight of 0.005. Making the overall loss equal to: DICE Loss $+ 0.005 \times \mathcal{L}_{topo}$ Further due to the computationally expensive nature of this loss, it was calculated on volumes resampled to $16 \times 16 \times 16$. Again, the same baseline neural network similar training parameters were used. Training results are shown below.

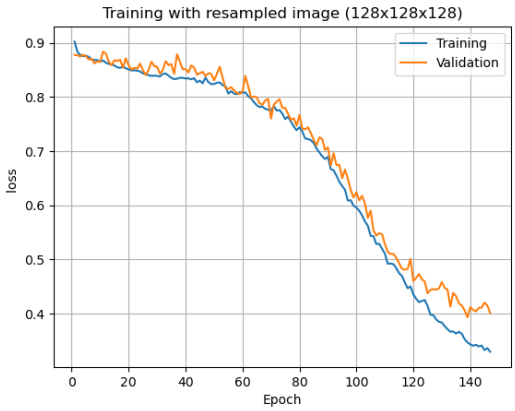

Figure 6.2: Training results with topological + DICE loss

Evident from the results, incorporating the topological loss didn't bring much improvement over the usual DICE loss.

Chapter 7

Interactive Segmentation

7.1 InterCNN

InterCNN is an interactive segmentation model proposed by Bredell *et al* [5]. The idea is train a CNN to update a given segmentation output, based on scribble inputs from a user. As such, there are two major parts to it: an automatic segmentation model (AutoCNN) that provides an initial, possibly less accurate segmentation, and another interactive part which learns to correct this initial segmentation based on user inputs (InterCNN). The following figure gives a general overview of the segmentation process:

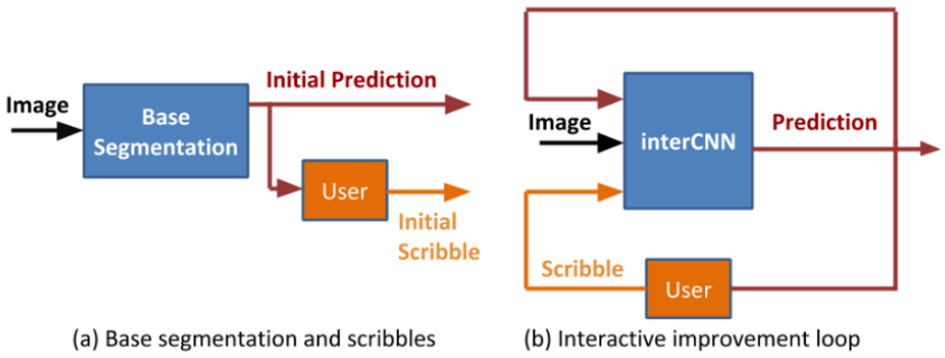

Figure 7.1: Segmentation process in InterCNN (*Source: [5]*)

Note that scribbles are supposed be a bunch of pixels on particular slices, marked by a user to be incorrectly classified. While training, both the automatic and interactive networks were chosen to be 2D UNets. The CT-scan volumes were split into slices and fed during training in batches of size 8. The Adam optimizer with a learning rate of 10^{-3} was used. Notably, the authors of this paper used the cross entropy loss during training and only used the DICE score

for validation. We used the same configuration.

Another important detail is that a robot user was to generate scribbles while training the interactive. This user randomly selected incorrectly classified pixels for each class in a given slice, and marked all the wrongly classified pixels around the selected pixel with a box of size 9×9 as the scribble. During training, for each slice, the number of interactions allowed with the robot user is fixed to be 10.

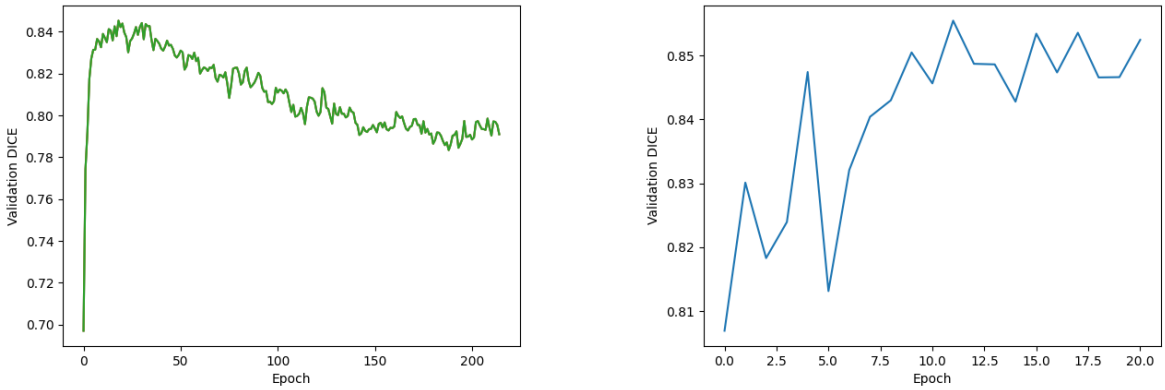

(a) Results for automatic part

(b) Results for interactive part

Figure 7.2: Training results for 2D InterCNN

The results show an overall increase in segmentation accuracy for the validation set with the incorporation of user interaction.

7.2 InterCNN-3D

While InterCNN produces much better results than automatic segmentation methods, it only works on a slice-by-slice basis, meaning that a user would have to interact with each slice of a volume individually in order to obtain a segmentation of a the whole volume. This is requires too much interaction, and calls for a method that requires less interaction, preferably with the whole volume instead of slices. Thus motivated, we tried to extend InterCNN to a 3D version, where a 3D network would learn to update a full 3D segmentation based on user inputs.

The segmentation networks from the original InterCNN can be easily translated to 3D (using 3D UNets instead of 2D). The only challenge is to decide on how a user would provide feedback to the interactive part, and creating a corresponding robot user for training. We tested multiple approaches to this, but haven't succeeded in producing results better than just automatic segmentation.

The first approach allows the user to draw scribbles on slices (along any of the three principle axes) of their choice in the volume, and then incorporate the scribbled slices back into a volume input. The second approach takes clicks on incorrect pixels from the user and produces volumetric Gaussian fields around those clicks in a volume. Results for incorporating both of these methods are shown in the following figures.

Figure 7.3: 3D AutoCNN (UNet) training

The AutoCNN results above are as expected for training of 3D UNet, with a validation DICE of 0.7. The following figures show training results for the other, interactive half, with different user-interaction policies.

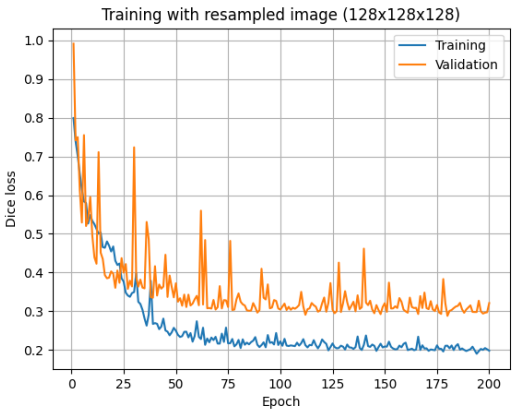

(a) With slice based interaction

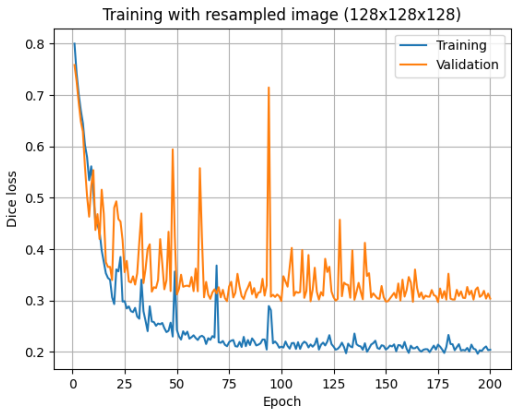

(b) With click based interaction

Figure 7.4: Training results for interactive part of InterCNN-3D with different interaction methods

Apart from the interaction, the training setup for both the automatic and interactive part

was similar to the how we trained 3D UNets before. We experimented with different number of interactions with the robot user during training. The results above correspond to training with 20 interactions. The results so far have failed to show any improvement upon the incorporation of user interaction. Possible causes maybe issues with the network architecture and its ability to detect and make use of user inputs, as well as the way interactions (scribbles/clicks) are encoded and passed to the model. This is still under investigation.

Chapter 8

Conclusion & Future Work

8.1 Conclusion

In conclusion, work in the project has so far focused on obtaining automatic or semi-automatic segmentation of CT volume data. We experimented with several techniques to try and improve the automatic segmentation results, after which we sought to utilize user interaction, which ended up yielding better results. There is still scope for reducing the number of interactions required to get a satisfactory segmentation, as well as experimenting with methods besides InterCNN for interactive segmentation. Coming back to aiding Brachytherapy treatments, future work will include developing a full-fledged and easy-to-use tool for interactive segmentation, followed by using the tool as a building block for development of a treatment simulation framework.

8.2 Future Work

Future work in the project consists of the following.

- Find and fix problems with InterCNN-3D.
- Build a fully-fledged GUI tool for segmentation.
- Incorporate obtained segmentations into patient specific treatment simulations.

Bibliography

[1] Brian M. Anderson, Kareem A. Wahid, and Kristy K. Brock. Simple Python Module for Conversions Between DICOM Images and Radiation Therapy Structures, Masks, and Prediction Arrays. *Practical Radiation Oncology*, 11(3):226–229, May 2021. ISSN 1879-8500. doi: 10.1016/j.prro.2021.02.003. 8

[2] Jeroen Bertels, Tom Eelbode, Maxim Berman, Dirk Vandermeulen, Frederik Maes, Raf Bisschops, and Matthew B. Blaschko. Optimizing the dice score and jaccard index for medical image segmentation: Theory and practice. *Medical Image Computing and Computer Assisted Intervention – MICCAI 2019*, page 92–100, 2019. ISSN 1611-3349. doi: 10.1007/978-3-030-32245-8_11. URL http://dx.doi.org/10.1007/978-3-030-32245-8_11. 10

[3] Y.Y. Boykov and M.-P. Jolly. Interactive graph cuts for optimal boundary and region segmentation of objects in n-d images. In *Proceedings Eighth IEEE International Conference on Computer Vision. ICCV 2001*, volume 1, pages 105–112 vol.1, 2001. doi: 10.1109/ICCV.2001.937505. 6

[4] Freddie Bray, Jacques Ferlay, Isabelle Soerjomataram, Rebecca L. Siegel, Lindsey A. Torre, and Ahmedin Jemal. Global cancer statistics 2018: GLOBOCAN estimates of incidence and mortality worldwide for 36 cancers in 185 countries. *CA: a cancer journal for clinicians*, 68(6):394–424, November 2018. ISSN 1542-4863. doi: 10.3322/caac.21492. 2

[5] Gustav Bredell, Christine Tanner, and Ender Konukoglu. Iterative interaction training for segmentation editing networks, 2018. URL https://arxiv.org/abs/1807.08555. v, 7, 16

[6] Abhishek Chatterjee, Surbhi Grover, Lavanya Gurram, Supriya Sastri, and Umesh Mahantshetty. Patterns of cervical cancer brachytherapy in India: Results of an online survey supported by the Indian Brachytherapy Society. *Journal of contemporary brachytherapy*, 11(6):527–533, December 2019. ISSN 2081-2841. doi: 10.5114/jcb.2019.90448.

[7] Özgün Çiçek, Ahmed Abdulkadir, Soeren S. Lienkamp, Thomas Brox, and Olaf Ronneberger. 3D U-Net: Learning Dense Volumetric Segmentation from Sparse Annotation. *arXiv:1606.06650 [cs]*, June 2016. v, 6, 9

[8] James Clough, Nicholas Byrne, Ilkay Oksuz, Veronika A. Zimmer, Julia A. Schnabel, and Andrew King. A topological loss function for deep-learning based image segmentation using persistent homology. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, pages 1–1, 2020. doi: 10.1109/TPAMI.2020.3013679. 14

[9] Antonio Criminisi, Toby Sharp, and Andrew Blake. GeoS: Geodesic Image Segmentation. In David Forsyth, Philip Torr, and Andrew Zisserman, editors, *Computer Vision – ECCV 2008*, Lecture Notes in Computer Science, pages 99–112, Berlin, Heidelberg, 2008. Springer. ISBN 978-3-540-88682-2. doi: 10.1007/978-3-540-88682-2_9. 6

[10] Kris Derks, Jacco Steenhuijsen, Hetty Berg, Saskia Houterman, Jeltsje Cnossen, Paul Haaren, and Katrien Jaeger. Impact of brachytherapy technique (2D versus 3D) on outcome following radiotherapy of cervical cancer. *Journal of Contemporary Brachytherapy*, 10, February 2018. doi: 10.5114/jcb.2018.73955.

[11] Harish Doraiswamy and Vijay Natarajan. Computing Reeb Graphs as a Union of Contour Trees. *IEEE transactions on visualization and computer graphics*, 19(2):249–262, February 2013. ISSN 1941-0506. doi: 10.1109/TVCG.2012.115.

[12] Marina Horn, Benjamin Reh, Frederik Wenz, Jan Stallkamp, and Katja Mombaur. Patient specific corotated FEM simulation and gelatin phantom for prostate brachytherapy. In *2016 6th IEEE International Conference on Biomedical Robotics and Biomechatronics (BioRob)*, pages 341–346, June 2016. doi: 10.1109/BIOROB.2016.7523649.

[13] Xiaoling Hu, Li Fuxin, Dimitris Samaras, and Chao Chen. Topology-Preserving Deep Image Segmentation. *arXiv:1906.05404 [cs]*, June 2019.

[14] Anuj Kumar, Supriya Chopra, Umesh Mahantshetty, and Vijay Anand Reddy. Brachytherapy in India: Learning from the past and looking into the future. *Brachytherapy*, 19(6): 861–873, November 2020. ISSN 1538-4721. doi: 10.1016/j.brachy.2020.08.019. 2

[15] Wenqi Li, Guotai Wang, Lucas Fidon, Sebastien Ourselin, M. Jorge Cardoso, and Tom Vercauteren. On the compactness, efficiency, and representation of 3d convolutional networks:

Brain parcellation as a pretext task. In *Lecture Notes in Computer Science*, pages 348–360. Springer International Publishing, 2017. doi: 10.1007/978-3-319-59050-9_28. URL https://doi.org/10.1007%2F978-3-319-59050-9_28. 6

[16] Zheng Lin, Zhao Zhang, Lin-Zhuo Chen, Ming-Ming Cheng, and Shao-Ping Lu. Interactive image segmentation with first click attention. In *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 13336–13345, 2020. doi: 10.1109/CVPR42600.2020.01335. 7

[17] Zheng Lin, Zhao Zhang, Lin-Zhuo Chen, Ming-Ming Cheng, and Shao-Ping Lu. Interactive Image Segmentation With First Click Attention. In *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 13336–13345, June 2020. doi: 10.1109/CVPR42600.2020.01335.

[18] Umesh Mahantshetty, Shivakumar Gudi, Roshni Singh, Ajay Sasidharan, Supriya (Chopra) Sastri, Lavanya Gurram, Dayanand Sharma, Selvaluxmy Ganeshrajah, Janaki MG, Dinesh Badakh, Abhishek Basu, Francis James, Jamema V Swamidas, Thayalan Kuppuswamy, and Rajendra Bhalavat. Indian Brachytherapy Society Guidelines for radiotherapeutic management of cervical cancer with special emphasis on high-dose-rate brachytherapy. *Journal of Contemporary Brachytherapy*, 11(4):293–306, August 2019. ISSN 1689-832X. doi: 10.5114/jcb.2019.87406.

[19] Fausto Milletari, Nassir Navab, and Seyed-Ahmad Ahmadi. V-net: Fully convolutional neural networks for volumetric medical image segmentation, 2016. URL https://arxiv.org/abs/1606.04797. 6

[20] A Panchal and R Keyes. SU-GG-T-260: Dicompyler: An Open Source Radiation Therapy Research Platform with a Plugin Architecture. *Medical Physics*, 37(6Part19):3245–3245, 2010. ISSN 2473-4209. doi: 10.1118/1.3468652. 8

[21] Fernando Pérez-García, Rachel Sparks, and Sébastien Ourselin. Torchio: a python library for efficient loading, preprocessing, augmentation and patch-based sampling of medical images in deep learning. *Computer Methods and Programs in Biomedicine*, page 106236, 2021. ISSN 0169-2607. doi: https://doi.org/10.1016/j.cmpb.2021.106236. URL https://www.sciencedirect.com/science/article/pii/S0169260721003102. 9

[22] Mathias Perslev, Erik Bjørnager Dam, Akshay Pai, and Christian Igel. One network to segment them all: A general, lightweight system for accurate 3d medical image segmentation. In *Lecture Notes in Computer Science*, pages 30–38. Springer International

Publishing, 2019. doi: 10.1007/978-3-030-32245-8_4. URL https://doi.org/10.1007%2F978-3-030-32245-8_4. 6

[23] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-Net: Convolutional Networks for Biomedical Image Segmentation. *arXiv:1505.04597 [cs]*, May 2015. 6

[24] Janusz Skowronek. Current status of brachytherapy in cancer treatment – short overview. *Journal of Contemporary Brachytherapy*, 9(6):581–589, December 2017. ISSN 1689-832X. doi: 10.5114/jcb.2017.72607.

[25] Carole H. Sudre, Wenqi Li, Tom Vercauteren, Sebastien Ourselin, and M. Jorge Cardoso. Generalised dice overlap as a deep learning loss function for highly unbalanced segmentations. In *Deep Learning in Medical Image Analysis and Multimodal Learning for Clinical Decision Support*, pages 240–248. Springer International Publishing, 2017. doi: 10.1007/978-3-319-67558-9_28. URL https://doi.org/10.1007%2F978-3-319-67558-9_28. 12

[26] Jack Yang. Oncentra brachytherapy planning system. *Medical Dosimetry: Official Journal of the American Association of Medical Dosimetrists*, 43(2):141–149, 2018. ISSN 1873-4022. doi: 10.1016/j.meddos.2018.02.011.