

**AI-based image auto-segmentation with AI-based accuracy assessment with respect to a  
clinical task: Application in Radiation Therapy**

A Technical Report submitted to the Department of Biomedical Engineering

Presented to the Faculty of the School of Engineering and Applied Science

University of Virginia • Charlottesville, Virginia

In Partial Fulfillment of the Requirements for the Degree

Bachelor of Science, School of Engineering

**Aaron Patton**

Spring, 2020.

Technical Project Team Members

Nicole Chomicki

On my honor as a University Student, I have neither given nor received unauthorized aid on this  
assignment as defined by the Honor Guidelines for Thesis-Related Assignments

Dr. Jeffery Siebers and Dr. Hussain Soomro, Department of MD-RONC Radiation Oncology

## **Abstract**

Head and neck cancer is estimated to be diagnosed in 66,630 Americans in 2021.<sup>1</sup> Radiation therapy (RT) is a treatment option that can be used to cure cancer or reduce its recurrence. The RT treatment planning process relies on physician expertise and currently involves manual segmentation (MS) and manual review (MR). This process, however, is time-consuming, laborious, and prone to error.<sup>2</sup> The project therefore aimed to alleviate these issues by creating a pipeline that automatically segments organs at risk (OARs), runs a Knowledge-Based Quality Control (KBQC) step on the segmentations, performs 3D dose prediction and estimates the dose metrics for each OAR and planning target volume (PTV), then determines which OARs are “relevant” and need to be further reviewed by the physician before final treatment planning begins. OARs labelled “irrelevant” have dose levels far below clinical tolerances and do not need further review. The pipeline was developed in a way that allows the user to easily run the pipeline with little to no coding experience. The project also aimed to quantify the usefulness of the pipeline in a clinical setting by analyzing the OAR relevancy prediction of the pipeline. The goal was that the pipeline had to have an average of at least one OAR labelled “irrelevant” so the physician would spend less time and labor in planning the RT treatment. The project used 33 publicly-available computed tomography (CT) scans from The Cancer Imaging Archive to perform these tests. The pipeline is able to analyze a CT scan in approximately 5 minutes. The pipeline had an average of 1.3 OARs deemed “irrelevant,” and 82% of patients had at least one OAR labelled “irrelevant.” Our results validate indicate that the pipeline should reduce time and labor for the physician and will be worth using in the future in a clinical setting.

Keywords: Head and Neck Cancer, Radiation Therapy, Organs at Risk, Pipeline, Automatic Segmentation, Relevancy

---

## **Introduction**

Head and neck cancer begins in several places in the head and throat, but does not include brain and eye cancers. The typical age of diagnosis is 50 years or older, and twice as many men as women are diagnosed with this disease. Alcohol and tobacco use are major risk factors for developing head and neck cancer.<sup>3</sup> It accounts for 4% of cancer patients in the United States, and it is estimated to kill approximately 14,620 Americans in the year 2021.<sup>1</sup> Patients currently have three main options for cancer treatment: radiation therapy (RT), chemotherapy, and surgery. Approximately one third of cancer patients elect to undergo RT.<sup>4</sup> RT uses intense radiation beams (most often X-rays) to kill cancer cells. A challenge in RT is to destroy the cancerous cells while minimizing collateral damage to nearby organs-at-risk (OARs). Therefore, RT aims to eliminate the cancerous tissue while saving surrounding OARs as possible.<sup>5</sup>

In order for RT to be effective, a treatment plan must be made to optimize where and how much radiation should be applied to destroy the tumor without damaging normal tissue. The basic RT process can be broken down into five steps. First, a CT or magnetic resonance image of the cancerous region must be acquired for the purpose of planning the patient’s treatment. Second, the image needs

to be segmented into target and surrounding OAR regions. Manual segmentation (MS) is typically used to delineate the target and OAR boundaries on the planning image. The third step is treatment planning. In this step, experts optimize the direction and intensity of the radiation beams to achieve a prescription dose to the target while keeping the dose to the OARs below complication inducing dose tolerances. The dose planning is an iterative process, which includes a manual review (MR) step in which the physician cross-checks the accuracy of each OAR segmentation and checks the tradeoffs between OAR and target doses. When target dose coverage is lacking, or OARs are receiving too much radiation, the treatment plan will be adjusted until clinically acceptable tradeoffs are achieved.<sup>6</sup> It is possible that planning loop is executed several times to develop an acceptable plan. With an acceptable plan, the patient will then undergo treatment.<sup>2</sup>

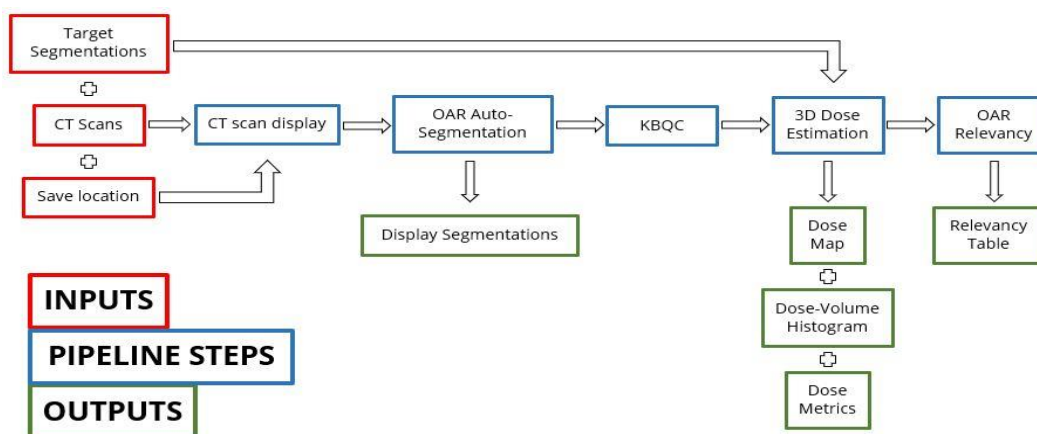
The above described RT workflow has several shortcomings. For example, MS is susceptible to multiple error modes and is inherently irreproducible at the millimeter level. MR is unlikely to find all of the relevant errors.<sup>2,7</sup> The standard RT workflow can be a long and laborious process – lasting from a minimum of several hours up to a maximum of more than a week when inherent idle times between steps is considered. If fewer OARs are segmented, the segmentation time can be reduced, however this is infrequently done since it is difficult to determine which OARs can be safely skipped. Skipping a relevant OAR could have substantial negative patient consequences. However, time spent on delineation and review of irrelevant OARs is wasted – time which could be better spent on other clinical activities. Therefore, there is a critical need for an automated workflow that can accurately predict the relevant OARs for each CT scan so that a efficient treatment workflow can be utilized. To achieve this task, a reliable, accurate, efficient and less error-prone segmentation method and dose prediction method are needed. Automatic segmentation and dose prediction methods using deep learning can be much faster than their manual counterpart and help achieve the desired automated workflow.<sup>8,9</sup>

To improve the RT workflow, we set out to create a pipeline that formalizes the flow of OAR segmentation, performs a knowledge based quality control (KBQC) step, estimates the optimized dose, then determines which OARs require physician review. The step-by-step flow of this pipeline is displayed in Figure 1. The KBQC library is meant to determine plausibility of the automatic segmentations and reduce gross errors. The novelty of our

pipeline is that all of the aforementioned steps are included in one interactive, easy-to-use user interface (UI). There is no current commercial software that implements this workflow, however, there are several RT treatment planning packages which are used in the typical RT workflow. At UVA, Pinnacle (Philips Medical Systems) is used. While Pinnacle does have auto-segmentation algorithms, it lacks any quality control step and users report that the time to correct the auto-segmentations exceeds the time it takes to manually segment the structures. Pinnacle has dose optimization tools, which can take several hours per iteration to complete.<sup>10</sup> As such, there is no quick way to pre-assess OAR relevancy. Our pipeline is meant to reduce RT segmentation errors, improve segmentation consistency, and reduce time and labor on the RT treatment planning process. Utilizing our novel pipeline can allow patients to receive RT treatments with less risk of damage to their OARs.

## Materials and Methods

Python was the coding language that was used to develop this pipeline. Jupyter Notebook, an open-source web application that can contain live code, equations, visualizations, and narrative text, was installed to make the backend of the code. Jupyter Notebook allowed us to easily separate the code into individual cells so that each section of the pipeline could be tested to ensure they functioned properly. Another Python package that we installed was Voilà. Voilà transforms Jupyter Notebooks into user-friendly standalone web applications. This was what was used for the frontend of our pipeline. For the relevancy analysis, the 33 datasets were acquired from The



**Fig. 1. Flowchart of the pipeline inputs, outputs, and step-by-step process.** The pipeline accepts the CT scans and folder path where the results will be saved as inputs. When the pipeline is run, the original CT scans are displayed and automatically segmented. The scans go through a KBQC process, then the doses of each OAR and PTV are calculated. Finally, Each OAR is analyzed by the pipeline to determine which ones are “relevant” and need to be further reviewed by the physician. The pipeline has several outputs, which are the segmentation results, dose map, dose-volume histogram, dose metrics, and a relevancy table.

Cancer Imaging Archive. 31 datasets were from Head-Neck\_Cetuximab.<sup>11</sup> 2 datasets were from Head-Neck-PET-CT.<sup>12</sup>

In order to run each step of the pipeline, the user must click buttons to select a folder or run a section of the pipeline. The Python packages that were used to create the buttons and attach them to particular functions were ipywidgets and tkinter. The tkinter package allows the user to select the folder that the CT scan files were located, as well as which folder the pipeline's results should be saved. The buttons and interactive sliders were developed by using ipywidgets.

To build the auto-segmentation step of the pipeline, ipywidgets was used again to make two buttons: the first button runs the library that automatically segments each OAR, and the second button displays the segmented CT scans. The library used to automatically segment each OAR is called "OARnet," which was developed by Dr. Soomro. OARnet is responsible for segmenting 28 OARs in the head and neck, and it is able to save each segmented scan as a .png file so the user can easily view the results.

Similar to the automatic segmentation step, the 3D Dose Estimation step was separated into two steps. The first step runs the dose estimation using a library called "DeepDoseNet." DeepDoseNet is capable of using the segmented targets and OARs to estimate a dose map. Then, the dose metrics can be computed and visually displayed using the dose map. A Python package called "plotly" was used to make the dose-volume histogram, bar plot, and table highly interactive. The library "DeepDoseNet" was also developed by Dr. Soomro.

The final step of the pipeline compares the doses for each OAR predicted by DeepDoseNet to the dose objectives that UVA uses for each OAR. The dose objectives are the maximum and mean doses in Grays that are appropriate for each OAR (Table S1). When this step is run, the code creates a table that determines which OARs are "relevant" or "irrelevant." Dr. Rashad developed the "Relevancy" library.

To evaluate whether our pipeline reduces time and labor, we ran an analysis to test if at least one OAR on average would be deemed "irrelevant." Equation 1 was used to compare the dose estimated by the pipeline with the dose objectives times a K value, which is a "safe limit."

In the analysis, 33 patients were evaluated, and the K-value used for the dose-comparison equation was 0.8. The 33 patient data sets were obtained from The Cancer Imaging Archive, which is publicly available. The 33 patients were run through the pipeline and the number of "relevant" OARs out of seven total OARs was counted for each CT scan. The mean number of "relevant" OARs was computed, as was the percentage of patients with at least one "irrelevant" OAR.

## **Results**

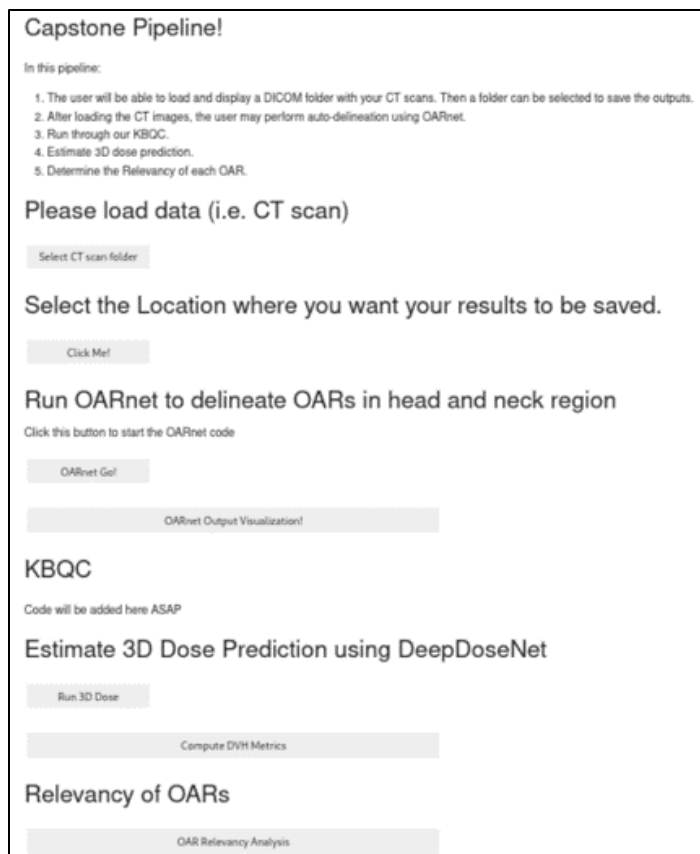
### ***Pipeline***

Our pipeline is currently able to display the original CT scan, perform auto-segmentation, undergo 3D dose estimation, and determine the relevancy of each OAR. When the user first opens the pipeline through Voila, that person is introduced to the main layout of the webpage. This layout is shown in Figure 2. The overall layout is a very simplistic design that utilizes interactive widgets and markdown code to separate each section of the pipeline.

The user is able to perform the first step of the pipeline by clicking a button and selecting the folder containing the CT images to be analyzed in DICOM format. When the user does this, the file path is stored as a variable to be used in the rest of the pipeline. The original CT scans are displayed and the user can scroll through each slice and change the image intensity via interactive sliders. A slice of the displayed CT scan images is shown in Figure 3A. The user can click another button to select in which folder the auto-segmentation results should be saved. Once the user selects the desired folder, the folder path will be printed on the pipeline. The user has the option to change the folder path if necessary.

$$D_{\text{estimated}} > K * D_{\text{objective}}$$

[1]

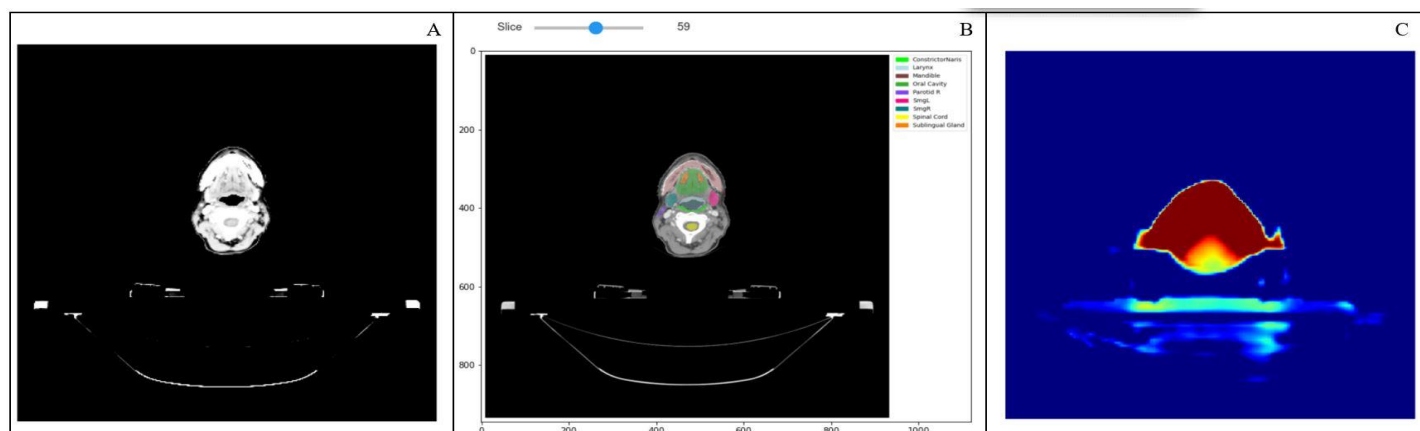


**Fig. 2. Overall layout of the pipeline’s frontend.** This is what users first see when they use Jupyter Notebook and Voila to open the pipeline. The pipeline has a simplistic overall layout and only requires the user to click buttons and select folders. Markdown is used to insert text into the webpage, and ipywidgets and tkinter are two Python packages that make the buttons function properly.

The auto-segmentation section has two buttons for the user to interact with. The first button begins the process for OAR segmentation. The OARnet library performs the auto-segmentation. This button also saves each segmented image slice as PNG files into the folder that the user had previously selected. Once this process is complete, the user then presses the second button, which displays the .png files. These images also have a legend on the side of each slice to explain where each OAR is located. One such slice is shown in Figure 3B.

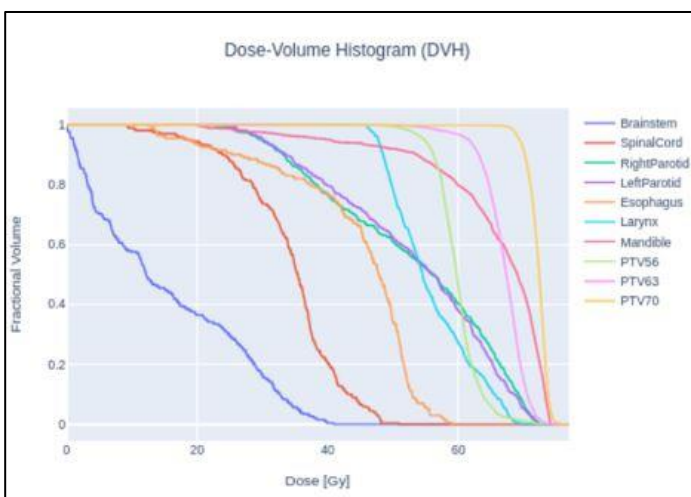
The KBQC has not yet been inserted into the pipeline because it is still being trained by Dr. Alves. The purpose of the KBQC is to find gross errors and establish which segmentations are plausible or absurd. The KBQC will work to reduce gross OAR segmentation errors. Once training is complete, it will be incorporated into the pipeline.

The 3D Dose Estimation section also contains two interactive buttons. When the first button is clicked, a dose map is computed and displayed. When the user scrolls the mouse over the image, the dose (in Grays, [Gy]) at the bottom of the image changes depending on the image slice and mouse location. This is displayed in Figure 3C. The user also has an intensity slide, similar to the slice shown in Figure 3A. The results from the dose map are then able to be used for the dose metrics calculations and relevancy code.



**Fig. 3. Displayed images of the original CT scan, OAR segmentation results, and dose map, respectively.** Each image is of the same slice. 3A shows the original CT scan, and the user has interactive sliders to change the slice number and image intensity. 3B shows the auto-segmentation results with a legend on the right side to show which color represents each OAR. The user has a slider to scroll through each segmented slice. 3C shows the dose map of the slice, and the displayed dose changes based on slice number and mouse location.


The first button in the 3D Dose Estimation section also produces a dose-volume histogram (DVH). A DVH is a visual representation of the dose within an OAR or PTV. It can help answer many questions about predicted dose based on segmented OARs and target volumes. For example, a physician may want to know what the percent volume is of a certain OAR that receives 40 Gy or more. The DVH in the pipeline is highly interactive. The user can click on the legend and hide certain OARs and PTVs that do not need to be seen. The user can also move their mouse over the plot and see the x and y values of one or all items in the legend, depending on which option the user wants. Figure 4 shows the DVH with all OARs and PTVs present on the plot.



**Fig. 4. DVH developed by the pipeline.** The DVH shows what fractional volume of each OAR and PTV would receive at various doses. The plot is interactive, as the user can hide and display any item in the legend. The user can select options where one or all coordinate values are shown when the mouse moves over a curve.

The second button in the 3D Dose Estimation section computes the dose metrics for each OAR and PTV. The results are displayed in a stacked bar chart and table format, shown in Figure 5 and Table 1. There are seven dose metrics that are computed for each OAR and PTV. These metrics are minimum dose, mean dose, maximum dose, D99%, D95%, D5%, and D1%. The last four metrics stand for what dose each OAR and PTV receives at the specified fractional volume. Figure 5 allows the user to compare how much dose one OAR or PTV receives relative to another OAR or PTV for each dose metric. For example, Figure 5 shows that the mandible will receive a higher mean dose than the brainstem. Because it is a stacked bar chart, its main purpose is to look at numeric values across two categorical variables.<sup>13</sup> Similar to Figure 4, Figure 5 is interactive. The items in its legend can be expressed or hidden, and the user can scroll their mouse over each metric to see the doses of one or all items. Table 1 takes the dose metric values from Figure 5 and relays that information in a tabular format. Each column in Table 1 can be clicked and dragged so that the user can organize each OAR and PTV to whichever layout is best for the physician.

The Relevancy section of the pipeline has one button. When the button is pressed, the estimated dose is compared to dose objectives that UVA Health System uses for plan optimization to determine which OARs are “relevant” and thus need to be further reviewed by the physician. The specific dose objectives for several OARs are shown in Table S1. The mean or maximum doses are compared, depending on which OAR is analyzed. An OAR is considered “relevant” if its estimated dose is larger than the dose objective times the K constant. This was displayed in Equation 1.



**Fig. 5. Stacked bar chart of the dose metrics.** The bar plot shows the max, min, and mean doses for each OAR and PTV. It also shows D99%, D95%, D5%, and D1%. These represent what dose each OAR and PTV receives at the specified fractional volume. This bar plot has the same interactive capabilities as the DVH in Fig. 4.

metrics	Brainstem	SpinalCord	RightParotid	LeftParotid	Esophagus	Larynx	Mandible	PTV56	PTV63	PTV70
Min[Gy]	0	9.3	22.4	23.8	12.6	45.9	16.7	38.6	48.3	63.7
Mean[Gy]	15.5	34.1	53	53.1	44	55.8	65.1	59.8	66.9	72.1
Max[Gy]	40.9	50.9	72.4	72.3	59	69.9	74.2	75.1	76.2	76.9
D99%[Gy]	0.2	10	25.3	25.9	13.1	46.3	21.7	51	55.1	67.8
D95%[Gy]	1	18.7	29.8	30.1	18.5	47.7	40.3	54.8	61.6	69.4
D5%[Gy]	35.5	46.4	70.4	70.2	55.3	67.3	73.6	65.2	71.3	74
D1%[Gy]	39.7	48.5	72.2	71.8	58.1	68.6	74	70	72.7	74.6

**Table 1. Tabular format of the dose metrics.** The table shows the max, min, mean, D99%, D95%, D5%, and D1% doses for each OAR and PTV. The table is interactive, as the user can click and drag each column to place them next to whichever other columns the user wants.

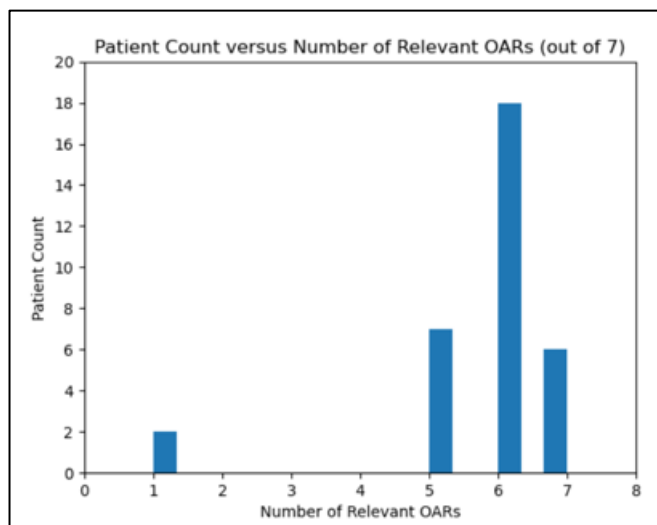
$D_{estimated}$  is the dose computed by the pipeline.  $D_{objective}$  is the dose objective that UVA Health System uses.  $K$  is a constant between 0 and 1. It is a “safe limit” that the pipeline sets. For our pipeline and relevancy analysis, 0.8 was the  $K$  value that was used. Another safety measure that the pipeline sets is that any OAR not listed in the dose objective list or segmented by OARnet is automatically deemed “relevant.” Table 2 shows the results that are produced after clicking the button on the pipeline. In this figure, all OARs for this patient are listed as “Relevant.”

	OAR	Relevancy
0	Brainstem	Relevant
1	SpinalCord	Relevant
2	RightParotid	Relevant
3	LeftParotid	Relevant
4	Esophagus	Relevant
5	Larynx	Relevant
6	Mandible	Relevant

**Table 2. Display of OARs analyzed for relevancy.** The final section of the pipeline determines and displays which analyzed OARs are “Relevant” (need to be reviewed by the physician) or “Irrelevant” (do not need to be reviewed by the physician). In this particular dataset, all OARs were deemed “Relevant.”

**Analysis**

After its successful creation, the pipeline was tested and analyzed to determine if it would be worth using in a clinical setting. 33 publicly available CT scan datasets were run through the pipeline. Of the 33 patients that were evaluated, 27 had at least 1 out of 7 OARs deemed “Irrelevant.” In other words, 82% of pipeline runs met the criteria for the pipeline to be worth using. 2 patients had only 1 “relevant” OAR. The average number of “relevant” OARs was 5.7 out of 7. Figure 6 shows a visual representation of the analysis results.



**Fig. 6. Analysis of patient count versus number of relevant OARs determined by the pipeline.** 33 publicly-available CT scans were evaluated, and the pipeline deemed an average of 1.3 OARs “Irrelevant” out of 7. 82% of datasets had at least 1 “Irrelevant” OAR.

**Discussion**

The results from the pipeline show that it is possible to create an automated, highly-interactive webpage that can aid in and enhance the RT treatment planning process. The pipeline has been shown to improve segmentation consistency and reduce time and labor on the planning process, as the pipeline takes less than five minutes to fully run and the segmentation results produce consistent results when datasets are run through the pipeline.

Based on the results from the analysis, the average number of “Irrelevant” OARs was 1.3. The criteria for our pipeline to be considered worth using in a clinical setting was if an average of at least 1 OAR or more was labelled “Irrelevant.” The results therefore met the goal, because having even one “irrelevant” OAR on average would save the physician time and labor. In this set of patients, there

was a relatively high average of “relevant” OARs. This is because the PTVs were located close to the analyzed OARs. Had the PTVs been located further away, the average number of “relevant” OARs would be lower. Seven potential OARs to head and neck radiation therapy were used in the analysis.

Our pipeline takes approximately five minutes to fully run. The auto-segmentation step takes the most time (about 2-3 minutes) because saving the segmented images as individual png files is a time-intensive task. However, this OARnet library is incredibly informative because it segments 28 OARs in the head and neck. Physicians often only segment OARs that have a close proximity to the PTV. The pipeline is therefore able to provide more information to the doctor and potentially reveal an unexpected “relevant” OAR. The main goal of the pipeline was to reduce time and labor for the physician by determining in a short amount of time which OAR(s), if any, are “irrelevant.” Based on our results, this goal was successfully completed.

### ***Limitations***

While our pipeline has many capabilities, there are limitations with our product. For example, the libraries of code were developed only for head and neck cancer CT scans. The reason for this is because head and neck cancer has less variability in organ position compared to other types of cancer. Also, the pipeline only accepts CT scans in DICOM file format. As previously mentioned in the Results section, our pipeline does not contain the KBQC yet. Our pipeline requires a high-performance computer GPU to be used. Otherwise, the pipeline cannot fully run.

Aside from limitations with the pipeline itself, the team was trying to determine a flexible and easy-to-use framework for establishing the pipeline. During the first half of the project, the team was attempting to use Django, a Python-based framework for webpages, to build the pipeline. However, progress on the pipeline was slower than desired when using Django, so the team switched to Jupyter Notebook and Voila for a simpler and more efficient building process. This hurdle set the team’s schedule back by two months. Progress was also occasionally hampered due to COVID-19 preventing the team from meeting in-person to quickly resolve issues.

### ***Future Steps***

In the future, the KBQC will be fully trained and will be inserted into the pipeline. The section of the pipeline that determines which OARs are relevant needs to be upgraded

to allow the user to easily alter the K constant based on what “safe limit” is desired. The code should also be upgraded so that the user can input other file types besides DICOM format into the pipeline. We also wish to add isodose lines into the dose map. In the future, the team also wants to perform more testing on the relevancy section of the code as well as the KBQC once it is finished. The extra testing is to ensure that each portion of the pipeline is as accurate and precise as possible. The layout of the pipeline also needs to be updated in the future. Currently, it is functional, but the layout’s form is very basic. The content can be adjusted to the center instead of to the left, and more color or graphics can be added to the pipeline. We also hope to alter the pipeline so that the entire process is run from one button click instead of several. Finally, we wish to expand this pipeline so that it can be utilized for other types of cancer, such as pancreatic, lung, and prostate cancer.

This project led to the development of a pipeline that automatically segments OARs, estimates the dose metrics of each item, and determines which OARs need to undergo MR by the physician. The short time it takes to run the pipeline, along with the ability of the pipeline to predict relevant and irrelevant OARs from 28 OARs, emphasizes that our pipeline will be worth using in a clinical setting. The creation of this pipeline and testing of the Relevancy section will allow us to develop a greater insight into applications of machine-learning in the medical field while improving the quality of care that cancer patients will receive during RT. Improving segmentation quality should reduce damage to OARs, decreasing the chance that patients will have negative effects on healthy tissue. This new and improved method for OAR assessment promises to be a valuable resource for the medical community and an opportunity to improve cancer patients’ quality of life.

### **End Matter**

#### ***Author Contributions and Notes***

Dr. Jeffrey Siebers designed the project and provided technical guidance, Nicole A. Chomicki and Aaron A. Patton built the pipeline and “CT scan display” section, Dr. M. Hussain Soomro wrote the “OARnet” and “DeepDoseNet” software, Dr. Victor Alves is training the KBQC, Dr. Hashir Rashad wrote the “Relevancy” code, Dr. Hashir Rashad, Dr. M. Hussain Soomro, Nicole A. Chomicki, and Aaron A. Patton analyzed data; and Nicole A. Chomicki and Aaron A. Patton wrote the paper. The authors declare no conflict of interest.



### ***Acknowledgments***

We would like to thank Dr. Siebers for his guidance and advising throughout the project. We would also like to thank Dr. Soomro for developing several of the pipeline's libraries, aiding in the analysis, and providing immensely helpful guidance throughout the project. Furthermore, we want to thank Dr. Rashad for developing the Relevancy code, helping us with coding issues, and aiding us in the analysis. We also wish to thank Dr. Alves for training the pipeline's KBQC code and helping us with coding & account issues.

### **References**

1. Head and Neck Cancer: Statistics. *Cancer.Net* (2021). Available at: <https://www.cancer.net/cancer-types/head-and-neck-cancer/statistics>. (Accessed: 1st May 2021).
2. Siebers, J. BME Capstone AI based image auto-segmentation with AI based accuracy assessment with respect to a clinical task Application in Radiation Therapy. (2020).
3. Head and neck cancers. *Centers for Disease Control and Prevention* (2020). Available at: <https://www.cdc.gov/cancer/headneck/index.htm>. (Accessed: 19th April 2021).
4. Siebers, J. Capstone Weekly Meeting. (2021).
5. Radiation therapy. *Mayo Clinic* (2020). Available at: <https://www.mayoclinic.org/tests-procedures/radiation-therapy/about/pac-20385162>. (Accessed: 26th April 2021).
6. Soomro, M. & Rashad, H. Capstone Weekly Meeting. (2021).
7. Gillebert, C. R., Humphreys, G. W. & Mantini, D. Automated delineation of stroke lesions using brain CT images. *NeuroImage: Clinical* **4**, 540–548 (2014).
8. Siebers, J. Clarification on Pipeline. (2020).
9. Zaidi, H. & Naqa, I. E. PET-guided delineation of radiation therapy treatment volumes: a survey of image segmentation techniques. *European Journal of Nuclear Medicine and Molecular Imaging* **37**, 2165–2187 (2010).
10. Radiation oncology treatment planning systems | Philips Healthcare. *Philips* <https://www.usa.philips.com/healthcares/solutions/radiation-oncology/radiation-treatment-planning>.
11. Kirby, J. Data From Head-Neck\_Cetuximab. *Cancer Imaging Archive* (2020). Available at: <https://wiki.cancerimagingarchive.net/display/Public/Head-Neck+Cetuximab>. (Accessed: 6th May 2021)
12. Nolan, T. & Camp, B. Head-Neck-PET-CT. *Cancer Imaging Archive* (2021). Available at: <https://wiki.cancerimagingarchive.net/display/Public/Head-Neck-PET-CT>. (Accessed: 6th May 2021)
13. Yi, M. A Complete Guide to Stacked Bar Charts. *Chartio* (2019). Available at: <https://chartio.com/learn/charts/stacked-bar-chart-complete-guide/>. (Accessed: 4th May 2021)

**Supplemental Material**

OAR	MaxDose (Gy)	MeanDose (Gy)
'Brainstem'	54	N/A
'Constrict_Musc'	N/A	50
'Eye_L'	45	38
'Eye_R'	45	38
'Larynx'	63	N/A
'Mandible'	75	N/A
'OpticChiasm'	44	N/A
'OpticNrv_L'	55	N/A
'OpticNrv_R'	55	N/A
'Parotid_L'	N/A	26
'Parotid_R'	N/A	26
'SpinalCord'	45	N/A
'Sublings_Glnd'	N/A	35
'Thyroid_Glnd'	N/A	N/A
'Trachea'	69	N/A

**Table S1. Dose Objectives for OARs.** The dose objectives shown in this supplemental table are the clinical dose tolerances for the tissues that UVA Health System uses. If the estimated dose for an OAR from the pipeline is less than its dose objective times the K value (0.8), then there is no need to review that OAR. This is a partial list of all the OAR dose objectives that the UVA Health System uses.