

Original Research Article

Autocontouring of primary lung lesions and nodal disease for radiotherapy based only on computed tomography images

Stephen Skett^a, Tina Patel^a, Didier Duprez^b, Sunnia Gupta^a, Tucker Netherton^c, Christoph Trauernicht^b, Sarah Aldridge^a, David Eaton^a, Carlos Cardenas^d, Laurence E. Court^{c,*}, Daniel Smith^a, Ajay Aggarwal^a

^a Guy's and St. Thomas' NHS Foundation Trust, London, United Kingdom

^b Stellenbosch University Faculty of Medicine and Health Sciences, Tygerberg Hospital, Cape Town, South Africa

^c The University of Texas MD Anderson Cancer Center, Houston, TX, United States

^d University of Alabama at Birmingham Hazelrigg-Salter Radiation Oncology Center, Birmingham, AL, United States



ARTICLE INFO

Keywords:

Auto-contouring
Lung disease
Radiotherapy
Computed tomography
Deep learning
GTV

ABSTRACT

Background and purpose: In many clinics, positron-emission tomography is unavailable and clinician time extremely limited. Here we describe a deep-learning model for autocontouring gross disease for patients undergoing palliative radiotherapy for primary lung lesions and/or hilar/mediastinal nodal disease, based only on computed tomography (CT) images.

Materials and methods: An autocontouring model (nnU-Net) was trained to contour gross disease in 379 cases (352 training, 27 test); 11 further test cases from an external centre were also included. Anchor-point-based post-processing was applied to remove extraneous autocontoured regions. The autocontours were evaluated quantitatively in terms of volume similarity (Dice similarity coefficient [DSC], surface Dice coefficient, 95th percentile Hausdorff distance [HD95], and mean surface distance), and scored for usability by two consultant oncologists. The magnitude of treatment margin needed to account for geometric discrepancies was also assessed.

Results: The anchor point process successfully removed all erroneous regions from the autocontoured disease, and identified two cases to be excluded from further analysis due to 'missed' disease. The average DSC and HD95 were 0.8 ± 0.1 and 10.5 ± 7.3 mm, respectively. A 10-mm uniform margin-distance applied to the autocontoured region was found to yield "full coverage" (sensitivity > 0.99) of the clinical contour for 64 % of cases. Ninety-seven percent of evaluated autocontours were scored by both clinicians as requiring no or minor edits.

Conclusions: Our autocontouring model was shown to produce clinically usable disease outlines, based on CT alone, for approximately two-thirds of patients undergoing lung radiotherapy. Further work is necessary to improve this before clinical implementation.

1. Introduction

Lung cancer is the leading cause of cancer deaths worldwide [1], and evidence-based estimates indicate that about three-quarters of all lung cancer patients would derive clinical benefit from treatment with radiotherapy to improve survival, reduce disease progression, or palliate the debilitating symptoms from progressive disease [2].

Computed tomography (CT) based radiotherapy requires that clinically relevant anatomic structures be contoured on the planning CT scan, including the gross tumour volume (GTV) and normal tissues at risk of significant radiation exposure. Contouring is a critical aspect of

radiotherapy treatment planning, which typically requires significant, time-consuming manual effort by clinicians and is subject to significant inter-user variabilities [3,4].

Automation of contouring has been identified as a solution that may address this challenge with respect to achieving more consistent, high-quality treatment and reducing the burden on clinical teams from these labour-intensive processes, enabling more efficient workforce allocation to other tasks in the radiotherapy pathways.

Although positron-emission tomography (PET) imaging can provide important information when contouring lung disease, this imaging modality is unavailable in many clinics across the world. There is a need,

* Corresponding author.

E-mail address: lecourt@mdanderson.org (L.E. Court).

<https://doi.org/10.1016/j.phro.2024.100637>

Received 4 April 2024; Received in revised form 21 August 2024; Accepted 21 August 2024

Available online 24 August 2024

2405-6316/© 2024 Published by Elsevier B.V. on behalf of European Society of Radiotherapy & Oncology. This is an open access article under the CC BY-NC-ND license (<http://creativecommons.org/licenses/by-nc-nd/4.0/>).

therefore, to develop autocontouring solutions which are based only on the radiotherapy planning CT scan. In this work specifically, we aimed to develop a deep-learning model for autocontouring gross lung disease. The principal target group for the model is patients undergoing palliative radiotherapy for primary lung lesions, with or without nodal involvement; if successful, the model could be considered for inclusion in a future version of the Radiation Planning Assistant automated treatment planning tool, which is being developed with the same ethos and goals in mind [5]. Although many researchers have described efforts to auto-contour lung tumors [6–10], we believe that ours is the first to focus on palliative radiotherapy using only CT images. Our long-term goal is to make automated contouring and planning available to clinical teams in low- and middle-income countries, helping them scale their efforts to treat more patients and potentially reducing the gaps in access to and quality of radiotherapy across the world.

2. Materials and methods

2.1. Datasets

A dataset was first created from patients treated at a single institution (Guy’s and St. Thomas’ NHS Foundation Trust, GSTT) from August 2016 through March 2021, and from October 2021 through October 2022. All data and images were reviewed by a dosimetrist or medical physicist, who consulted on a case-by-case basis with other members of the clinical team, including a thoracic oncology registrar (resident) and specialist consultants, to establish the final cohort using the inclusion/exclusion criteria stated in Table 1.

This procedure resulted in a final dataset of 379 cases (231 radical, 148 palliative), which was randomly split by using a readily available Python-based pseudo-random number generator into 352 training/cross-validation : 27 testing. Most of the testing cases (15 of 27) were primary tumor only (T0) and most of the other cases (9 of 12) had ipsilateral nodal disease (N1/2). The final testing set was supplemented with 11 additional cases from a second institution (Tygerberg Hospital, South Africa) according to the same inclusion/exclusion criteria,

Table 1
Inclusion and exclusion criteria for the training/testing dataset.

| Inclusion criteria | |
|------------------------------|--|
| Clinical intent | Treatment of lung cancer (or other thoracic malignancy involving the lung, e.g., lymphoma), including associated hilar/mediastinal malignant lymphadenopathy |
| Imaging | Planning CT scan with clinician-outlined CT-visible gross disease: this included lung tumor(s) and any involved lymph nodes When valid clinical contours were available for any of a set of standard thoracic organs at risk (OARs; i.e., lungs, heart, oesophagus and spinal cord), these were also included as inputs during training |
| Exclusion criteria | |
| Disease location | Treatment of supraclavicular fossa / neck nodes |
| Treatment approach | Receipt of adjuvant radiotherapy after surgery |
| Target visibility | Some portion of treated disease was identified solely via positron emission tomography (PET) or endobronchial ultrasound and outlined on the basis of these findings Lesion edge could not be visually distinguished from adjacent non-lung tissues or lung abnormalities (e.g., such as collapse-consolidation or atelectasis) Tumor maximum dimension < 1 cm |
| Target selection | Additional tumour and/or nodal disease region(s) present and visible on the scan but excluded from the treatment plan outlining (i.e., cases in which radiotherapy was not intended to cover all gross disease, or specific disease regions were treated with separate plans & CT scans) |
| Incomplete target contouring | Some/all treated disease was not contoured during radiotherapy planning (e.g., when treatment involved simple field placement and accurate disease outlining was not undertaken) |

providing a total of 38 patients as a final testing set. All data extraction and analyses were approved by the ethics committee at the treating institution.

2.2. Autocontouring

An autocontouring model was trained with the nnU-Net architecture (v 1.7) [11], using the default values of the nnU-Net-determined rules-based hyperparameters. This architecture was selected because of its high level of success reported in the literature [12,13]. The input data for training was the CT images, original clinical GTV contours, and standard organs at risk (lungs, heart, esophagus and spinal cord). Images were resized, windowed, and z-normalized during pre-processing using parameters automatically determined by the nnU-Net framework. Training began with the He initialization and stochastic gradient descent using Nesterov momentum [14]. We used nnU-Net’s default hyperparameters [11]: 5-fold cross-validation, initial learning rate 0.01, combined Dice and cross-entropy loss, SGD optimizer with Nesterov momentum ($\mu = 0.99$), batch size 2, and 1000 epochs. Data augmentation included rotations, scaling, noise, blur, brightness/contrast adjustments, and mirroring. The framework automatically optimized patch size, network depth, and features per layer.

During early development of the model, the autocontouring models were found to identify volumes that are distinct from those that the physician originally treated, particularly in the setting of multifocal bilateral disease. Because the eventual goal of this work is to develop a single-step, end-to-end workflow that includes contouring and radiotherapy treatment planning without the need for human intervention, the potential for mistaken identification of disease regions is a significant risk factor with severe consequences (i.e., treating the wrong volume). Therefore, we introduced the use of target-region anchor points, as illustrated in Fig. 1. These points would be used by the clinician at the pre-planning stage to identify regions that they wish to treat. Any autogenerated volumes that do not intersect with an anchor point may then be automatically removed. This approach gives the physician full control over which regions are auto-contoured for treatment. For the purposes of generating data for this current study, anchor points were placed by a medical physicist, in approximately the center of the original clinical volumes (that is, the clinical volumes were used by the physicist to ensure that the anchor points were placed in the region(s) that the original physician intended to treat).

2.3. Quantitative evaluation

Automatically generated contours were evaluated by comparing them with the contours created by clinicians (“clinical contours”) using four common geometric and volumetric similarity metrics: Dice similarity coefficient (DSC), 95th percentile Hausdorff distance (HD95), mean surface distance (MSD), and surface Dice coefficient (SDC) with 2-mm tolerance [15–18].

Uncertainties are commonly accounted for during radiotherapy planning by expanding the outlined structure by a certain margin. Although these expansions are generally determined based on setup and motion, they also account for contouring variations. The use of margins was investigated by applying a series of uniform margins to each auto-contour, ranging from 2 mm to 30 mm. The expanded structure was then compared with the clinical volume, and pixel-wise sensitivity was calculated. The expanded structure was considered to ‘fully cover’ the clinical structure when sensitivity > 0.99 was achieved.

2.4. Subjective evaluation

All test cases were anonymized and reviewed independently by two consultant clinical oncologists and scored on a 4-point scale: (1) unusable, (2) major edits necessary, (3) minor edits necessary, (4) use as-is (no edits necessary). The reviewers were provided with radiological

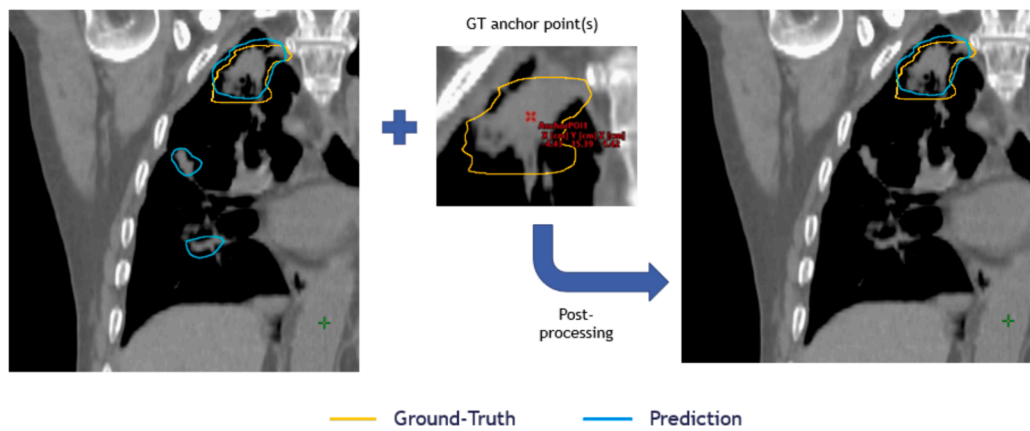


Fig. 1. Illustration of the use of anchor points (GT: Ground Truth).

staging information and the positions of the pre-planning anchor-points, but not the original clinical target outlines or any information from the diagnostic imaging.

3. Results

3.1. Use of anchor points

The autocontouring model identified a total of 77 discrete regions across the 38 test cases. Of these, 29 regions were correctly removed by the use of anchor points; the remaining 48 autocontoured regions were all correctly-identified discrete lesions or nodal regions. In 2 of the 38 test cases (5%), the model did not contour a volume that encompassed an anchor point—that is, it contoured a completely different region not contiguous with the intended treatment location. Those two cases were excluded from further quantitative analysis, on the basis that the presence of ‘missed’ anchor points in the predicted GTV could be easily detected and excluded as part of built-in quality-control self-checks in any future deployment of this model.

Over the remaining 36 test cases, the autocontoured outline for 3 of the 48 volumes across 3 separate test cases (8%) were found to represent only a partial visual match to the overall clinical contour (based on shape and location), in such a way that the ability of anchor points to accept/reject an autocontour was judged to be limited and may depend on the exact location of the anchor point in question. That is, the medical physicist subjectively determined that a manually-placed anchor point would likely be positioned outside the automatically identified structure.

3.2. Quantitative evaluation

The results of the geometric evaluation are illustrated in Fig. 2. The




Fig. 2. Box plots showing the results of the quantitative evaluation of the autocontours.

average DSC was 0.78 ± 0.10 and HD95 was 10.5 ± 7.3 mm. The box plots show close-to-normal distributions with relatively little skew for both the DSC and SDC, but with a significantly greater spread in the SDC than in the DSC. This indicates that variations at the level of surface changes around the contour edge are more pronounced, compared with the relatively high level of overall volumetric agreement between the clinical and autocontoured outlines.

Conversely, the distributions of HD95 and MSD showed significant positive skew, with the mean being greater than the median for both. This was predominantly due to a small number of outlier cases that had particularly large discrepancy distances: for example, if the two most extreme anomalous cases are discounted, the adjusted average HD95 would be 9.3 ± 5.5 mm.

When data were separated by origin (GSTT and Tygerberg), the average DSC was 0.80 ± 0.11 and 0.76 ± 0.07 , respectively. Similarly, SDC was 0.59 ± 0.20 and 0.43 ± 0.10 ; HD95 was 8.6 ± 5.7 mm and 14.7 ± 9.3 mm; and MSD was 2.6 ± 1.5 mm and 4.0 ± 1.9 mm. These data indicate that the model performs best on the same patient population as was used for training (i.e., GSTT data) but that reasonable performance is maintained when tested on an outside dataset (i.e., Tygerberg data).

Fig. 3 shows the effect of adding different uniform margins around the autocontour. For 23 of the 36 cases analysed (64%), a 10-mm margin was sufficient to cover the clinical contour; while an asymmetric margin at 10 mm in the lateral and anterior-posterior (AP) directions and 15 mm in the superior-inferior direction would cover 25 of 36 cases (69%), as seen in Fig. 3. Two illustrative exemplar cases are also shown: the first (A) is a typical case in which the standard planning margin would lead to adequate target coverage; the other (B) is an example where the autocontour does not match the clinical contour well. Notably, the purpose of this exercise was to assess whether margins could potentially account for variations in the autocontours, and does not consider any of the other factors that are included when determining clinical planning margins (for example, patient motion or setup uncertainties).

3.3. Subjective evaluation

Twenty-one of the cases reviewed (58%) were scored as 4 (no edits) by both reviewers. Thirty-five (97%) were scored either 3 (minor edits) or 4 (no edits). The single remaining case was scored 1 (unusable). The level of overall agreement between the two reviewers was very high, with 30 of the 36 cases (83%) receiving the same score. The DSC and HD95 for the failed case was 0.7 and 34 mm, respectively; although DSC was within the same range as found for all other cases, HD95 was significantly outside that range.

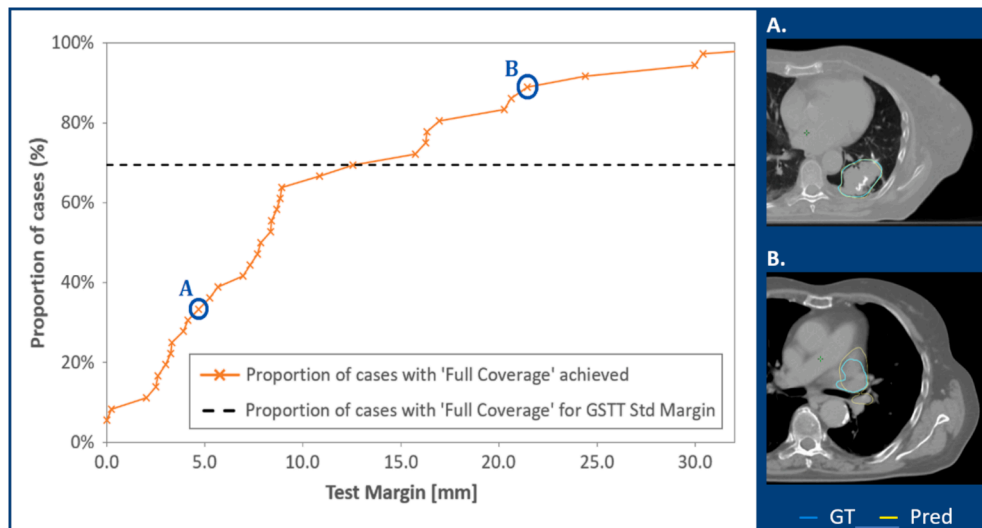


Fig. 3. The percentage of test cases for which the autocontour, expanded by a test margin, covers the clinical contour, as a function of margin distance. The right panel shows two example cases.

4. Discussion

This work showed that a commonly-used deep learning architecture can be applied to the autocontouring of radiotherapy targets for the treatment of lung cancer, using CT images alone. Of the 36 test cases reviewed (after filtering out cases in which known regions of disease were ‘missed’ by the autocontoured outlines), 97 % were subjectively scored by our clinicians as requiring no more than minor edits. The cases used were representative of those typically treated in clinical practice, including 11 from South Africa (an independent data set).

The validity of using pre-planning anchor points as a simple and robust method for removing unintended regions from the auto-generated target contours was also demonstrated: the anchor-points–based post-processing successfully removed all incorrectly identified disease regions from the autocontoured outlines across the test set, as well as providing a means of quickly identifying and excluding cases in which the autocontouring had ‘missed’ true disease regions. This is particularly important if the autocontouring model is to be incorporated into an autoplanning workflow, because incorrect contouring could result in treatment of incorrect volumes. The use of PET imaging may reduce the need for this human input, but PET is not available for many cancer patients, especially in limited-resource settings or when the goal of the treatment is palliation. Three cases were manually identified for which the anchor point would have likely not resulted in the correct identification of the target structure. Two of these were relative small/narrow shaped targets (nodal disease), where it was thought that small variations in the location of the anchor point could affect the success of the use of the anchor point. Quantitative assessment for these cases found that the geometric metrics (DSC, etc.) were not outliers, and could not guide these subjective assessments. For these cases, there is simply a high likelihood that the autocontours would be incorrectly rejected (based on anchor placement). Additional approaches for using anchor points are also possible and may give different results, including cropping the image around an anchor point before training/inference or providing the anchors in a separate channel in the nnU-Net.

Many reports have been published on autocontouring for lung cancer planning, most of which focus on OARs (lungs, esophagus, spinal cord, heart, etc.) and generally show good agreement between autogenerated and manual contours [3,19,20]. Published work on the autocontouring of lung tumors generally shows DSC values similar to those reported here. For example, Gan et al [7] reported an average DSC of 0.72, Ferrante et al [6] reported 0.78, and Hosny et al [8] reported 0.83, although

variations in datasets and study methodologies make direct comparisons challenging. Other researchers report similar, or worse, results [9], although again variations in datasets and study methodologies make direct comparisons challenging. Others have taken advantage of additional imaging modalities, such as PET, to give impressive DSC (0.83) and clinical acceptability [10], but we focused on CT because many clinics are unlikely to have access to these modalities.

This work also provided evidence that a standard treatment margin may be used to account for contouring uncertainties such that about two-thirds of cases could be planned without edits. Treatment planning margins are commonly used in radiotherapy to account for uncertainties and motion of the target. These include an internal margin to account for variations in location of the target due to, for example, respiratory motion, and an additional margin to account for other uncertainties such as contouring and target alignment. Systematic uncertainties (such as those from contouring) tend to dominate the size of the necessary margin. In short, determining the appropriate margin is complicated and is beyond the scope of this study, which did not consider most of these uncertainties. Thus, the margin presented above is not the final treatment margin, but serves more as an illustration of how uncertainties in the autocontour (and potentially the clinical contour) can be accounted for by using margins, and the required magnitude of such margins. Further work is necessary to determine how to safely apply this approach to an automated contouring/planning workflow. The degree to which margins can or should be used will also depend on the treatment approach. Simple techniques such as AP: posterior-anterior (PA) are more robust to variability/uncertainty in target contouring than are highly conformal techniques such as intensity-modulated radiation therapy or volumetric modulated arc therapy, especially for hypofractionated treatment.

An important aspect of our findings was that the subjective review was not just unambiguously positive, but also consistent between our two clinician reviewers. Although others have noted that expert segmentation style and preference may affect clinical utility for lung tumor contouring [8] and even for normal tissue contouring [21], we found a high level of agreement between the two consultants. Overall, nearly 60 % of the autocontours were scored as acceptable without any edits required by both of our oncologist reviewers. The quantitative results were similar in that about 70 % of cases would have achieved full target dose coverage if a realistic treatment margin were applied.

Although testing was performed using CT datasets from two different institutions, this does not necessarily mean the results would be maintained at other locations, where CT image quality and/or patient

populations may be different. Our experience suggests that CT image quality will not be the major driver for any differences, partly because operational CT scanners tend not to generate poor images and partly because deep learning-based autocontouring is not overly sensitive to the CT parameters, as long as they are kept to a reasonable range. Patient populations, however, can vary widely, and we have less experience with this. As is true for all clinical tools, extensive testing on a wide patient population will be necessary before clinical deployment.

Whilst these findings do suggest that the remaining cases would need some modification before use, the time savings may still be substantial. In similar work involving fully automated radiotherapy for post-mastectomy breast treatments, edits were found to be needed for about 50 % of cases, but those edits could be made in 12–13 min, which is much shorter than the 2 h needed for a treatment planner to generate a plan from scratch [22]. We can expect similar results for lung planning using our model.

In order to use this tool clinically, it must first be integrated into an appropriate interface, such as a commercial treatment planning system, or the Radiation Planning Assistant which has been developed specifically to support clinics with limited resources [5,23], followed by an implementation study to establish the optimal way to ensure that the tool is used clinically. Of particular interest is whether the autocontouring quality is sufficient to be beneficial in clinical use.

In conclusion, our autocontouring model was shown to produce clinically usable disease outlines, based on CT alone, for approximately two-thirds of lung radiotherapy cases. Although this model is not sufficient for use in an independent end-to-end autocontouring process, it may be applicable as a starting point for automated planning, provided that adequate manual oversight from suitably-trained clinicians is included in the overall process.

Disclosures

Work for this study was funded by the Wellcome Trust.

Declaration of Competing Interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

References

- [1] Sung H, Ferlay J, Siegel RL, Laversanne M, Soerjomataram I, Jemal A, et al. Global Cancer Statistics 2020: GLOBOCAN estimates of incidence and mortality worldwide for 36 cancers in 185 countries. *CA Cancer J Clin* 2021;71:209–49. <https://doi.org/10.3322/caac.21660>.
- [2] Vinod SK, Hau E. Radiotherapy treatment for lung cancer: Current status and future directions. *Respirol* 2020;25:61–71. <https://doi.org/10.1111/resp.13870>.
- [3] Lustberg T, van Soest J, Gooding M, Peressutti D, Aljabar P, van der Stoep J, et al. Clinical evaluation of atlas and deep learning based automatic contouring for lung cancer. *Radiother Oncol* 2018;126:312–7. <https://doi.org/10.1016/j.radonc.2017.11.012>.
- [4] Cardenas CE, Yang J, Anderson BM, Court LE, Brock KB. Advances in auto-segmentation. *Semin Radiat Oncol* 2019;29:185–97. <https://doi.org/10.1016/j.semradonc.2019.02.001>.
- [5] Court L, Aggarwal A, Burger H, Cardenas C, Chung C, Douglas R, et al. Addressing the global expertise gap in radiation oncology: the radiation planning assistant. *JCO Glob Oncol* 2023:e2200431. Doi: 10.1200/go.22.00431.
- [6] Ferrante M, Rinaldi L, Botta F, Hu X, Dolp A, Minotti M, et al. Application of nnU-Net for automatic segmentation of lung lesions on CT images and its implication for radiomic models. *J Clin Med* 2022;11:7334. <https://doi.org/10.3390/jcm11247334>.
- [7] Gan W, Wang H, Gu H, Duan Y, Shao Y, Chen H, et al. Automatic segmentation of lung tumors on CT images based on a 2D & 3D hybrid convolutional neural network. *Brit J Radiol* 2021;94:20210038. <https://doi.org/10.1259/bjr.20210038>.
- [8] Hosny A, Bitterman DS, Guthrie CV, Qian JM, Roberts H, Perni S, et al. Clinical validation of deep learning algorithms for radiotherapy targeting of non-small-cell lung cancer: an observational study. *Lancet Digital Health* 2022;4:e657–66. [https://doi.org/10.1016/S2589-7500\(22\)00129-7](https://doi.org/10.1016/S2589-7500(22)00129-7).
- [9] Liu X, Li K-W, Yang R, Geng L-S. Review of deep learning based automatic segmentation for lung cancer radiotherapy. *Front Oncol* 2021;11. <https://doi.org/10.3389/fonc.2021.717039>.
- [10] Wang S, Mahon R, Weiss E, Jan N, Taylor RJ, McDonagh PR, et al. Automated lung cancer segmentation using a PET and CT Dual-modality deep learning neural network. *Int J Radiat Oncol Biol Phys* 2023;115:529–39. <https://doi.org/10.1016/j.ijrobp.2022.07.2312>.
- [11] Isensee F, Jaeger PF, Kohl SAA, Petersen J, Maier-Hein KH. nnU-Net: a self-configuring method for deep learning-based biomedical image segmentation. *Nat Methods* 2021;18:203–11. <https://doi.org/10.1038/s41592-020-01008-z>.
- [12] Wasserthal J, Breit H-C, Meyer MT, Pradella M, Hinck D, Sauter AW, et al. Totalsegmentator: robust segmentation of 104 anatomic structures in CT Images. *Radiol Artif Intell* 2023;5. <https://doi.org/10.1148/ryai.230024>. e230024.
- [13] Yu C, Anakwenze CP, Zhao Y, Martin RM, Ludmir EBS, Niedzielski J, et al. Multi-organ segmentation of abdominal structures from non-contrast and contrast enhanced CT images. *Sci Rep* 2022;12. <https://doi.org/10.1038/s41598-022-21206-3>. 19093.
- [14] Bottou L. In: *Large-Scale Machine Learning with Stochastic Gradient Descent*. Heidelberg: Physica-Verlag HD; 2010. p. 177–86.
- [15] Sorensen T. A method of establishing groups of equal amplitude in plant sociology based on similarity of species content and its application to analyses of the vegetation on Danish commons. *Biologiske skrifter* 1948;5:1–34.
- [16] Heimann T, Van Ginneken B, Styner MA, Arzhaeva Y, Aurich V, Bauer C, et al. Comparison and evaluation of methods for liver segmentation from CT datasets. *IEEE Trans Med Imaging* 2009;28:1251–65. <https://doi.org/10.1109/TMI.2009.2013851>.
- [17] Nikolov S, Blackwell S, Zverovitch A, Mendes R, Livne M, De Fauw J, et al. Clinically applicable segmentation of head and neck anatomy for radiotherapy: deep learning algorithm development and validation study. *J Med Internet Res* 2021;23:e26151. <https://doi.org/10.2196/26151>.
- [18] Vaassen F, Hazelaar C, Vaniqui A, Gooding M, Van der Heyden B, Canters R, et al. Evaluation of measures for assessing time-saving of automatic organ-at-risk segmentation in radiotherapy. *Phys Imaging Radiat Oncol* 2020;13:1–6. <https://doi.org/10.1016/j.phro.2019.12.001>.
- [19] Vaassen F, Boukerroui D, Looney P, Canters R, Verhoeven K, Peeters S, et al. Real-world analysis of manual editing of deep learning contouring in the thorax region. *Phys Imaging Radiat Oncol* 2022;22:104–10. <https://doi.org/10.1016/j.phro.2022.04.008>.
- [20] Walls GM, Giacometti V, Apte A, Thor M, McCann C, Hanna GG, et al. Validation of an established deep learning auto-segmentation tool for cardiac substructures in 4D radiotherapy planning scans. *Phys Imaging Radiat Oncol* 2022;23:118–26. <https://doi.org/10.1016/j.phro.2022.07.003>.
- [21] Baroudi H, Brock KK, Cao W, Chen X, Chung C, Court LE, et al. Automated contouring and planning in radiation therapy: what is clinically acceptable? *Diagnostics* 2023;13:667. <https://doi.org/10.3390/diagnostics13040667>.
- [22] Esho T, Chung C, Thompson J, Dehghanpour M, Sutton J, Shaitelman S, et al. Optimization of autogenerated chest-wall radiation treatment plans developed for postmastectomy breast cancer patients in underserved clinics. *Med Dosim* 2020;45. <https://doi.org/10.1016/j.meddos.2019.12.003>.
- [23] Court LE. The Radiation Planning Assistant: addressing the global gap in radiotherapy services. *Lancet Oncol* 2024;25:277–8. [https://doi.org/10.1016/S1470-2045\(24\)00084-6](https://doi.org/10.1016/S1470-2045(24)00084-6).