



Deep Esophageal Clinical Target Volume Delineation Using Encoded 3D Spatial Context of Tumors, Lymph Nodes, and Organs At Risk

Dakai Jin^{1(✉)}, Dazhou Guo¹, Tsung-Ying Ho^{2(✉)}, Adam P. Harrison¹,
Jing Xiao³, Chen-kan Tseng², and Le Lu¹

¹ PAII Inc., Bethesda, MD, USA
jindakai376@paii-labs.com

² Chang Gung Memorial Hospital, Linkou, Taiwan, ROC
tyho@cgmh.org.tw

³ Ping An Technology, Shenzhen, China

Abstract. Clinical target volume (CTV) delineation from radiotherapy computed tomography (RTCT) images is used to define the treatment areas containing the gross tumor volume (GTV) and/or sub-clinical malignant disease for radiotherapy (RT). High intra- and inter-user variability makes this a particularly difficult task for esophageal cancer. This motivates automated solutions, which is the aim of our work. Because CTV delineation is highly context-dependent—it must encompass the GTV and regional lymph nodes (LNs) while also avoiding excessive exposure to the organs at risk (OARs)—we formulate it as a deep contextual appearance-based problem using encoded spatial contexts of these anatomical structures. This allows the deep network to better learn from and emulate the margin- and appearance-based delineation performed by human physicians. Additionally, we develop domain-specific data augmentation to inject robustness to our system. Finally, we show that a simple 3D progressive holistically nested network (PHNN), which avoids computationally heavy decoding paths while still aggregating features at different levels of context, can outperform more complicated networks. Cross-validated experiments on a dataset of 135 esophageal cancer patients demonstrate that our encoded spatial context approach can produce concrete performance improvements, with an average Dice score of $83.9 \pm 5.4\%$ and an average surface distance of 4.2 ± 2.7 mm, representing improvements of 3.8% and 2.4 mm, respectively, over the state-of-the-art approach.

1 Introduction

Esophageal cancer ranks the sixth in global cancer mortality [1]. As it is usually diagnosed at rather late stage [18], radiotherapy (RT) is a cornerstone of treatment. Delineating the 3D clinical target volume (CTV) on a radiotherapy computed tomography (RTCT) scan is a key challenge in RT planning. As

Fig. 1 illustrates, the CTV should spatially encompass, with a mixture of pre-defined and judgment-based margins, primary tumor(s), *i.e.*, the gross tumor volume (GTV), regional lymph nodes (LNs) and sub-clinical disease regions, while simultaneously limiting radiation exposure to organs at risk (OARs) [2].




Fig. 1. Esophageal cancer CTV delineation, where red, yellow, and cyan indicate the GTV, regional LNs and CTV, respectively. (a) shows that the CTV is not a uniform margin expansion (brown-dotted line) from the GTV, while (b)–(d) shows how delineation becomes more complicated when regional LNs are present. (c) and (d) also depict wide and long examples of esophageal CTV, respectively. (Color figure online)

Esophageal clinical target volume (CTV) delineation is uniquely challenging because tumors may potentially spread along the entire esophagus and metastasize up to the neck or down to the upper abdomen LNs. Current clinical protocols rely on manual CTV delineation, which is very time and labor consuming and is subject to high inter- and intra-observer variability [12]. This motivates automated approaches to the CTV delineation.

Deep convolutional neural networks (CNNs) have achieved notable successes in segmenting semantic objects, such as organs and tumors, in medical imaging [4, 6–10]. However, to the best of our knowledge, no prior work, CNN-based or not, has addressed esophageal cancer CTV segmentation. Works on CTV segmentation of other cancer types mostly operate based on the RTCT appearance alone [14, 15]. As shown in Fig. 1, CTV delineation depends on the radiation oncologist’s visual judgment of both the appearance *and* the spatial configuration of the GTV, LNs, and OARs, suggesting that only considering the RTCT makes the problem ill-posed. Supporting this, Cardenas *et al.* recently showed that considering the GTV and LN binary masks together with the RTCT can boost oropharyngeal CTV delineation performance [3]. However, the OARs were not considered in their work. Moreover, binary masks do not explicitly provide distances to the model. Yet CTV delineation is highly driven by distance-based margins to other anatomical structures of interest, and it is difficult to see how regular CNNs could capture these precise distance relationships with binary masks alone.

Our work fills this gap by introducing a spatial-context encoded deep CTV delineation framework. Instead of expecting the CNN to learn distance-based margins from the GTV, LN, and OAR binary masks, we provide the CTV delineation network with the 3D signed distance transform maps (SDMs) [16] of these structures. Specifically, we include the SDMs of the GTV, LNs, lung, heart, and spinal canal with the original RTCT volume as inputs to the network. From a clinical perspective, this allows the CNN to emulate the oncologist’s manual delineation, which uses the distances of GTV and LNs vs. the OARs as a key constraint in determining CTV boundaries. To improve robustness, we randomly choose manually and automatically generated organ at risk OAR SDMs during training, while augmenting the GTV and LNs SDMs with the domain-specific jittering. We adopt a 3D progressive holistically nested network (PHNN) [6] to serve as our delineation model, which enjoys the benefits of strong abstraction capacities and multi-scale feature fusion with a light-weighted decoding path. We extensively evaluate our approach using a 3-fold cross-validated dataset of 135 esophageal cancer patients. Since we are the first to tackle automated esophageal cancer CTV delineation, we compare against previous CTV delineation methods for other cancers [3, 15], using the 3D PHNN as the delineation model. When comparing against pure appearance-based [15] and binary-mask-based [3] solutions, we show that our approach provides improvements of 10% and 3.8% in Dice score, respectively, with analogous improvements in Hausdorff distance (HD) and average surface distance (ASD). Moreover, we also show that PHNN is responsible for providing improvements of 1% in Dice score and 0.4 mm reduction in ASD over a 3D U-Net model [4].

2 Methods

CTV delineation in RT planning is essentially a margin expansion process, starting from observable tumorous regions (GTV and regional LNs) and extending into the neighboring regions by considering the possible tumor spread margins and distances to nearby healthy OARs. Figure 2 depicts an overview of our method, which consists of four major modularized components: (1) segmentation of prerequisite regions; (2) SDM computation; (3) domain-specific data augmentation; and (4) a 3D PHNN to execute the CTV delineation.

2.1 Prerequisite Region Segmentation

To provide spatial context/distance of the anatomical structures of interest, we must first know their boundaries. We assume that manual segmentations for the esophageal GTV and regional LNs are available. However, we do not make this assumption for the OARs. Indeed, missing organ at risk (OAR) segmentations ($\sim 20\%$) is common in our dataset. For the OARs, we consider three major organs: the lung, heart, and spinal canal, since most esophageal CTVs are closely integrated with these organs. Using the available organ labels, we trained a 2D

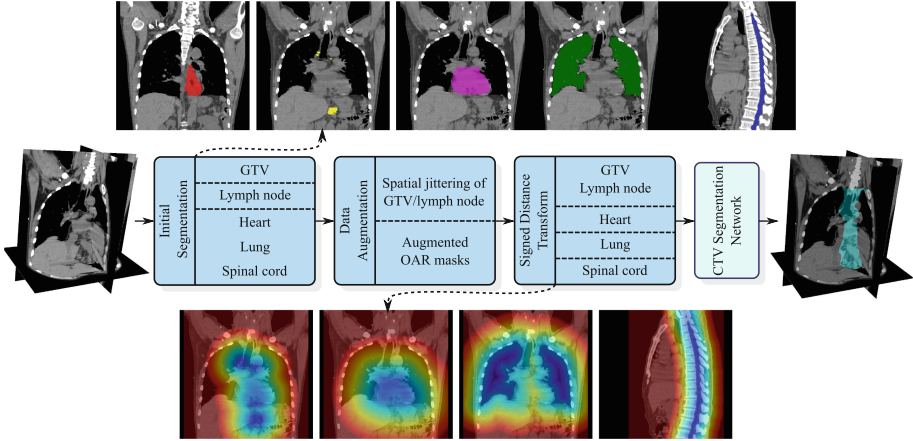


Fig. 2. Overall workflow of our spatial context encoded CTV delineation framework. The top and bottom rows depict different masks and SDMs, respectively, overlaid on the RTCT. From left to right are the GTV, LNs, heart, lung, and spinal canal. The GTV and LNs share a combined SDM.

PHNN [6] to segment the OARs, considering its robust performance in pathological lung segmentation and its computational efficiency. Examples of automatic OAR segmentation are illustrated in the first row in Fig. 2 and validation Dice score for the lung, heart and spinal canal were 97%, 95% and 78%, respectively, in our dataset.

2.2 SDM Computation

To encode the spatial context with respect to the GTV, regional LNs, and OARs, we compute signed distance transform maps (SDMs) for each. The SDM is generated from a binary image, where the value in each voxel measures the distance to the closest object boundary. Voxels inside and outside the boundary have positive and negative values, respectively. More formally, let \mathcal{O}_i denote a binary mask, where $i \in \{\text{GTV+LNs, lung, heart, spinal canal}\}$ and let $F(\cdot)$ be a function that computes boundary voxels of a binary image. The SDM value at a voxel p with respect to \mathcal{O}_i is computed as

$$\text{SDM}_{F(\mathcal{O}_i)}(p) = \begin{cases} \min_{q \in F(\mathcal{O}_i)} d(p, q) & \text{if } p \notin \mathcal{O}_i \\ -\min_{q \in F(\mathcal{O}_i)} d(p, q) & \text{if } p \in \mathcal{O}_i \end{cases}, \quad (1)$$

where $d(p, q)$ is a distance measure from p to q . We choose to use Euclidean distance in our work and use Maurer *et al.*'s efficient algorithm [13] to compute the SDMs. The bottom row in Fig. 2 depicts example SDMs for the combined GTV and LNs and the other 3 OARs. Note that we compute SDMs separately

for each of the three OARs, meaning we can capture each organ’s influence on the CTV. Providing the SDMs of the GTV, LNs, and OARs to the deep convolutional neural network (CNN) allows it to more easily infer the distance-based margins to these anatomical structures, better emulating the oncologist’s CTV inference process.

2.3 Domain-Specific Data Augmentation

We adopt specialized data augmentations to increase the robustness of the training and harden our network to noise in the prerequisite segmentations. Specifically, two types of data augmentation are carried out. (1) We calculate the GTV and LNs SDMs from both the manual annotations and also spatially jittered versions of those annotations. We jitter each GTV and lymph node (LN) component by random shift within $4 \times 4 \times 4 \text{ mm}^3$, mimicking that in practice 4 mm average distance error represents the state-of-the-art performance in esophageal GTV segmentation [8, 17]. (2) We calculate SDMs of the OARs using both the manual annotations and the automatic segmentations from Sect. 2.1. Combined, these augmentations lead to four possible combinations, which we randomly choose between during every training epoch. This increases model robustness and also allows the system to be effectively deployed in practice by using SDMs of the automatically segmented OARs, helping to alleviate the labor involved.

2.4 CTV Delineation Network

To use 3D CNNs in medical imaging, one has to strike a balance between choosing the appropriate image size covering enough context and the GPU memory. The symmetric encoder-decoder segmentation networks, *e.g.*, 3D U-Net [4], are computationally heavy and memory-consuming since half of its computation is consumed on the decoding path, which may not always be needed for all 3D segmentation tasks. To alleviate the computational/memory burden, we adopt a 3D version of PHNN [6] as our CTV delineation network, which is able to fuse different levels of features using parameter-less deep supervision. We keep the first 4 convolutional blocks and adapt it to 3D as our network structure. As we demonstrate in the experiments, the 3D PHNN is not only able to achieve reasonable improvement over the 3D U-Net but requires 3 times less GPU memory.

3 Experiments and Results

To evaluate the performance of our esophageal CTV delineation framework, we collected from 135 anonymized RTCTs of esophageal cancer patients undergoing RT. Each RTCT is accompanied by a CTV mask annotated by an experienced oncologist, based on a previously segmented GTV, regional LNs, and OARs. The average RTCT size is $512 \times 512 \times 250$ voxels with the average resolution of $1.05 \times 1.05 \times 2.6$ mm.

Training Data Sampling: We first resample all the CT and SDM images to a fixed resolution of $1.0 \times 1.0 \times 2.5$ mm, from which we extract $96 \times 96 \times 64$ training volume of interest (VOI) patches in two manners: (1) To ensure enough VOIs with positive CTV content, we randomly extract VOIs centered within the CTV mask. (2) To obtain sufficient negative examples, we randomly sample ~ 20 VOIs from the whole volume. This results in on average 80 VOIs per patient. We further augment the training data by applying random rotations of $\pm 10^\circ$ in the x-y plane.

Implementation Details: The Adam solver [11] is used to optimize all segmentation models with a momentum of 0.99 and a weight decay of 0.005 for 30 epochs. We use the Dice loss for training. For testing, we use 3D sliding windows with sub-volumes of $96 \times 96 \times 64$ and strides of $64 \times 64 \times 32$ voxels. The probability maps of sub-volumes are aggregated to obtain the whole volume prediction taking on average 6–7s to process one input volume using a Titan-V GPU.

Comparison Setup and Metrics: We use 3-fold cross-validation, separated at the patient level, to evaluate performance of our approach and the competitor methods. We compare against setups using only the CT appearance information [14,15] and setups using the CT with binary GTV/LN masks [3]. Finally, we also compare against setups using the CT + GTV/LN SDMs, which does not consider the OARs. We compare these setups using the 3D PHNN. For the 3D U-Net [4], we compared against the setup using the computed tomography (CT) appearance information. We evaluate the performance using the metrics of Dice score, ASD and HD.

Table 1. Quantitative results for the esophageal cancer CTV delineation.

Models	Setups	Dice	HD (mm)	ASD (mm)
U-Net	CT	0.739±0.126	69.5±42.7	10.1±9.4
	CT + GTV/LN/OAR SDMs	0.829±0.061	36.9±23.8	4.6±3.0
PHNN	CT	0.739±0.117	68.5±43.8	10.6±9.2
	CT + GTV/LN masks	0.801±0.075	56.3±35.4	6.6±5.3
	CT + GTV/LN SDMs	0.816±0.067	44.7±25.1	5.4±4.1
	CT + GTV/LN/OAR SDMs	0.839±0.054	35.4±23.7	4.2±2.7
	CT + GTV/LN/OAR SDMs*	0.823±0.059	43.6±26.4	5.1±3.3

*The last, starred row represents performance when using automatically generated OAR SDMs.

Results: Table 1 outlines the quantitative comparisons of the different model setups and choices. As can be seen, methods based on pure CT appearance, seen in prior art [14,15], exhibits the worst performance. This is because inferring distance-based margins from appearance alone is too hard of a task for CNNs. Focusing on the PHNN performance, when adding the binary GTV and LN




Fig. 3. Qualitative illustration of esophageal CTV delineation using different PHNN setups. Red, yellow and cyan represent the GTV, LN and predicted CTV regions, respectively. The purple line indicates the ground truth CTV boundary. The 1st and 2nd rows show examples from setups using pure RTCT [15] and when adding GTV/LN binary masks [3], respectively. The 3rd and 4th row show examples when adding GTV/LN SDMs and our proposed GTV/LN/OAR SDMs, respectively. (a) and (d) demonstrate that the pure RTCT setups fail to include the regional LNs, while (c) to (e) depict severe over-segmentations. While these errors are partially addressed using the GTV/LN mask setup, it still suffers from inaccurate CTV boundaries (a–c) or over coverage of normal regions (d, e). These issues are much better addressed by our proposed method. (Color figure online)

masks as contextual information [3], the performance increases considerably from 0.739 ± 0.117 to 0.801 ± 0.075 in Dice score. When using the SDM encoded spatial context of GTV/LN, PHNN further improves the Dice score and ASD by 1.5% and 1.2 mm, respectively, confirming the value of using the distance information for esophageal CTV delineation. Finally, when the OAR SDMs are included, *i.e.*, our proposed framework, PHNN achieves the best performance reaching 0.839 ± 0.054 Dice score and 4.2 ± 2.7 mm ASD, with a reduction of 9.3 mm in HD as compared to the next best PHNN result. Figure 4 depicts cumulative histograms of the Dice score and ASD, visually illustrating the distribution of improvements in the CTV delineation performance. Figure 3 shows some qualitative examples

illustrating these performance improvements. Interestingly, as the last row of Table 1 shows, when using SDMs computed from the automatically segmented OARs for testing, the performance compares favorably to the best configuration, and outperforms all other configurations. This indicates that our method remains robust to noise within the OAR SDMs and also that our approach is not reliant on manual OAR masks for good performance, increasing its practical value.




Fig. 4. Cumulative histograms of the CTV delineation performance under 4 setups using 3D PHNN on cross-validated 135 patients. The left and right depict the Dice score and ASD results, respectively. From the results, we observe that $> 77\%$ patients have Dice score ≥ 0.80 , and $> 55\%$ patients have Dice score ≥ 0.85 by using the proposed method (shown in red). Since there are often large inter-observer variations on CVT delineation tasks, *i.e.*, ranging from 0.51 to 0.81 in terms of Jaccard index in cervix cancer [5], these findings may indicate that, for a high percentage of the studied patient population, little to no additional manual revision is needed on the automatically delineated CTVs. (Color figure online)

We also compare the 3D PHNN network performance with that of 3D U-Net [4] when using the CT appearance based setup and the proposed whole framework. As Table 1 demonstrates, when using the whole pipeline PHNN outperforms U-Net by 1% dice score. Although PHNN has similar performance against U-Net when using only the CT appearance information, the GPU memory consumption is roughly 3 times less than that of the U-Net. These results indicate that for esophageal CTV delineation, a CNN equipped with strong encoding capacity and a light-weight decoding path can be as good as (or even superior to) a heavier network with a symmetric decoding path.

4 Conclusion

We introduced a spatial-context encoded deep esophageal CTV delineation framework designed to produce superior margin-based CTV boundaries. Our system encodes spatial context by computing the SDMs of the GTV, LNs and OARs and feeds them together with the RTCT image into a 3D deep CNN. Analogous to clinical practice, this allows the system to consider both appearance and distance-based information for delineation. Additionally, we also developed domain-specific data augmentation and adopted a 3D PHNN to further improve

robustness. Using extensive three-fold cross-validation, we demonstrated that our spatial-context encoded approach can outperform state-of-the-art CTV alternatives by wide margins in Dice score, HD, and ASD. As we are the first to address automated esophageal CTV delineation, our method represents an important step forward for this important problem.

References

1. Bray, F., Ferlay, J., et al.: Global cancer statistics 2018: Globocan estimates of incidence and mortality worldwide for 36 cancers in 185 countries. *CA a Cancer J. Clin.* **68**(6), 394–424 (2018)
2. Burnet, N.G., Thomas, S.J., Burton, K.E., Jefferies, S.J.: Defining the tumour and target volumes for radiotherapy. *Cancer Imaging* **4**(2), 153 (2004)
3. Cardenas, C.E., Anderson, B.M., et al.: Auto-delineation of oropharyngeal clinical target volumes using 3D convolutional neural networks. *Phys. Med. Biol.* **63**(21), 215026 (2018)
4. Çiçek, Ö., Abdulkadir, A., Lienkamp, S.S., Brox, T., Ronneberger, O.: 3D U-Net: learning dense volumetric segmentation from sparse annotation. In: Ourselin, S., Joskowicz, L., Sabuncu, M.R., Unal, G., Wells, W. (eds.) *MICCAI 2016*. LNCS, vol. 9901, pp. 424–432. Springer, Cham (2016). https://doi.org/10.1007/978-3-319-46723-8_49
5. Eminowicz, G., McCormack, M.: Variability of clinical target volume delineation for definitive radiotherapy in cervix cancer. *Radiother. Oncol.* **117**(3), 542–547 (2015)
6. Harrison, A.P., Xu, Z., George, K., Lu, L., Summers, R.M., Mollura, D.J.: Progressive and multi-path holistically nested neural networks for pathological lung segmentation from CT images. In: Descoteaux, M., Maier-Hein, L., Franz, A., Janin, P., Collins, D.L., Duchesne, S. (eds.) *MICCAI 2017*. LNCS, vol. 10435, pp. 621–629. Springer, Cham (2017). https://doi.org/10.1007/978-3-319-66179-7_71
7. Heinrich, M.P., Oktay, O., Bouteldja, N.: OBELISK-Net: fewer layers to solve 3D multi-organ segmentation with sparse deformable convolutions. *Med. Image Anal.* **54**, 1–9 (2019)
8. Jin, D., Guo, D., Ho, T.Y., et al.: Accurate esophageal gross tumor volume segmentation in pet/ct using two-stream chained 3d deep network fusion. In: Shen, D., et al. (eds.) *MICCAI 2019*. LNCS, vol. 11765, pp. 182–191. Springer, Cham (2019)
9. Jin, D., Xu, Z., Harrison, A.P., George, K., Mollura, D.J.: 3D convolutional neural networks with graph refinement for airway segmentation using incomplete data labels. In: Wang, Q., Shi, Y., Suk, H.-I., Suzuki, K. (eds.) *MLMI 2017*. LNCS, vol. 10541, pp. 141–149. Springer, Cham (2017). https://doi.org/10.1007/978-3-319-67389-9_17
10. Jin, D., Xu, Z., Tang, Y., Harrison, A.P., Mollura, D.J.: CT-realistic lung nodule simulation from 3D conditional generative adversarial networks for robust lung segmentation. In: Frangi, A.F., Schnabel, J.A., Davatzikos, C., Alberola-López, C., Fichtinger, G. (eds.) *MICCAI 2018*. LNCS, vol. 11071, pp. 732–740. Springer, Cham (2018). https://doi.org/10.1007/978-3-030-00934-2_81
11. Kingma, D.P., Ba, J.: Adam: a method for stochastic optimization. [arXiv:1412.6980](https://arxiv.org/abs/1412.6980) (2014)
12. Louie, A.V., Rodrigues, G., et al.: Inter-observer and intra-observer reliability for lung cancer target volume delineation in the 4D-CT era. *Radiother. Oncol.* **95**(2), 166–171 (2010)

13. Maurer Jr., C.R., Qi, R., Raghavan, V.: A linear time algorithm for computing exact euclidean distance transforms of binary images in arbitrary dimensions. *IEEE Trans. Pattern Anal. Mach. Intell.* **25**(2), 265–270 (2003)
14. Men, K., Dai, J., Li, Y.: Automatic segmentation of the clinical target volume and organs at risk in the planning ct for rectal cancer using deep dilated convolutional neural networks. *Med. Phys.* **44**(12), 6377–6389 (2017)
15. Men, K., Zhang, T., et al.: Fully automatic and robust segmentation of the clinical target volume for radiotherapy of breast cancer using big data and deep learning. *Physica Med.* **50**, 13–19 (2018)
16. Sethian, J.: A fast marching level set method for monotonically advancing fronts. *Proc. Natl. Acad. Sci.* **93**(4), 1591–1595 (1996)
17. Yousefi, S., et al.: Esophageal gross tumor volume segmentation using a 3D convolutional neural network. In: Frangi, A.F., Schnabel, J.A., Davatzikos, C., Alberola-López, C., Fichtinger, G. (eds.) *MICCAI 2018*. LNCS, vol. 11073, pp. 343–351. Springer, Cham (2018). https://doi.org/10.1007/978-3-030-00937-3_40
18. Zhang, Y.: Epidemiology of esophageal cancer. *World J. Gastroenterol. WJG* **19**(34), 5598 (2013)