



Lymph Node Gross Tumor Volume Detection and Segmentation via Distance-Based Gating Using 3D CT/PET Imaging in Radiotherapy

Zhuotun Zhu^{1,2(✉)}, Dakai Jin¹, Ke Yan¹, Tsung-Ying Ho³, Xianghua Ye⁵, Dazhou Guo¹, Chun-Hung Chao⁴, Jing Xiao⁶, Alan Yuille², and Le Lu¹

¹ PAII Inc., Bethesda, MD, USA
zhuotun@gmail.com

² Johns Hopkins University, Baltimore, MD, USA

³ Chang Gung Memorial Hospital, Linkou, Taiwan, ROC

⁴ National Tsing Hua University, Hsinchu City, Taiwan, ROC

⁵ The First Affiliated Hospital Zhejiang University, Hangzhou, China

⁶ Ping An Technology, Shenzhen, China

Abstract. Finding, identifying and segmenting suspicious cancer metastasized lymph nodes from 3D multi-modality imaging is a clinical task of paramount importance. In radiotherapy, they are referred to as Lymph Node Gross Tumor Volume (GTV_{LN}). Determining and delineating the spread of GTV_{LN} is essential in defining the corresponding resection and irradiating regions for the downstream workflows of surgical resection and radiotherapy of various cancers. In this work, we propose an effective distance-based gating approach to simulate and simplify the high-level reasoning protocols conducted by radiation oncologists, in a divide-and-conquer manner. GTV_{LN} is divided into two subgroups of “tumor-proximal” and “tumor-distal”, respectively, by means of binary or soft distance gating. This is motivated by the observation that each category can have distinct though overlapping distributions of appearance, size and other LN characteristics. A novel multi-branch detection-by-segmentation network is trained with each branch specializing on learning one GTV_{LN} category features, and outputs from multi-branch are fused in inference. The proposed method is evaluated on an in-house dataset of 141 esophageal cancer patients with both PET and CT imaging modalities. Our results validate significant improvements on the mean recall from 72.5% to 78.2%, as compared to previous state-of-the-art work. The highest achieved GTV_{LN} recall of 82.5% at 20% precision is clinically relevant and valuable since human observers tend to have low sensitivity ($\sim 80\%$ for the most experienced radiation oncologists, as reported by literature [5]).

Keyword: Lymph Node Gross Tumor Volume (GTV_{LN}), CT/PET Imaging, 3D Distance Transformation, Distance-based Gating

1 Introduction

Assessing the lymph node (LN) status in oncology clinical workflows is an indispensable step for the precision cancer diagnosis and treatment planning, e.g., radiation therapy or surgical resection. The class of enlarged LN is defined by the revised RECIST guideline [15] if its short axial axis is more than 10-15 mm in computed tomography (CT). In radiotherapy treatment, both the primary tumor and all metastasis suspicious LNs must be sufficiently treated within the clinical target volume with the proper doses [7]. We refer these LNs as lymph node gross tumor volume or GTV_{LN} , which includes enlarged LNs, as well as smaller ones that are associated with a high positron emission tomography (PET) signal or any metastasis signs in CT [14]. Accurately identifying and delineating GTV_{LN} , to be spatially included in the treatment area, is essential for a desirable cancer treatment outcome [10].

It is an extremely challenging and time-consuming task to identify GTV_{LN} , even for experienced radiation oncologists. High-level sophisticated clinical reasoning guidelines are needed, leading to the risk of uncertainty and subjectivity with high inter-observer variabilities [5]. It is arguably more difficult than detecting the more general enlarged LNs. (1) Finding GTV_{LN} is often performed using radiotherapy CT (RTCT) that (unlike diagnostic CT) is not contrast-enhanced. Hence the metastasis signs for identifying GTV_{LN} are subtler. (2) GTV_{LN} itself has poor contrast. Because of the shape and appearance ambiguity, it can be easily confused with vessels or muscles. (3) The size and shape of GTV_{LN} vary considerably with large amounts of smaller ones that are harder to detect. Refer Fig. 1 (top row) for an illustration of GTV_{LN} . While many previous works attempt to detect enlarged LNs using contrast-enhanced CT [1, 2, 4, 11–13, 18], no work, as of yet, has studied the GTV_{LN} detection in non-contrast RTCT scans. Given the evident differences between the enlarged LNs and GTV_{LN} , further innovations are required for the robust GTV_{LN} detection and segmentation.

Valuable insights from physicians’ clinical diagnosis and analysis process can be leveraged to tackle this problem. As one of the primary cues, human observers condition the analysis of GTV_{LN} based on the LNs’ distance with respect to the corresponding primary tumor location. For LNs proximal to the tumor, physicians more readily identify them as GTV_{LN} in radiotherapy treatment. However, for LNs distal to the tumor, they use more strict criteria to include if there are clear signs of metastasis, e.g., enlarged size, increased PET signals, and/or other CT based evidence [14]. Hence, the distance measure relative to the primary tumor plays a key role during physician’s decision making. Besides the distance, the PET modality is also of high importance. Although as a noisy imaging channel, it has shown to be helpful in increasing the GTV_{LN} detection sensitivity [5]. As demonstrated in Fig. 1 (bottom row), PET provides critically distinct information, yet, it also exhibits false positives (FPs) and false negatives (FNs).

In this paper, we imitate the physician’s diagnosis process to tackle the problem of GTV_{LN} detection and segmentation. (1) We introduce a distance-based gating strategy in a multi-task framework to divide the underlying GTV_{LN} distributions into “tumor-proximal” and “tumor-distal” categories and solve them

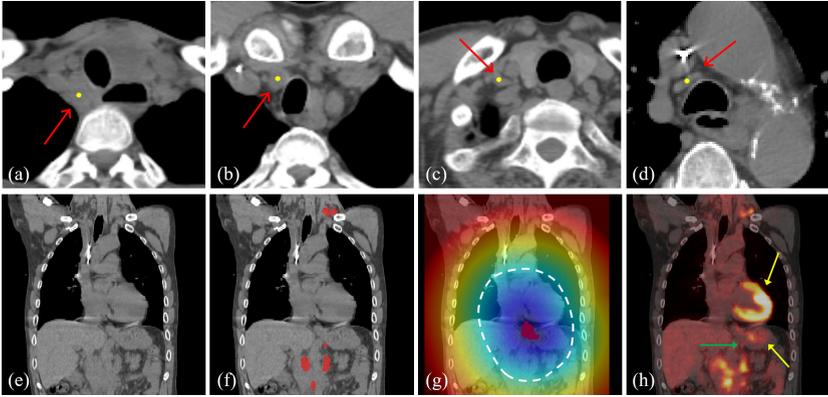


Fig. 1. Top row (a-d): examples of the GTV_{LN} (red arrow) with varying size and appearance at scatteredly distributed locations. Bottom row (e-h): (e) A coronal view of RTCT for an esophageal cancer patient. (f) The manual annotated GTV_{LN} mask. (g) The tumor distance transformation map overlaid on RTCT, where the primary tumor is indicated by red in the center and the white dash line shows an example of the binary tumor proximal and distal region division. (h) PET imaging shows several FPs with high signals (yellow arrows). Two FN GTV_{LN} are indicated by green arrow where PET has even no signals on a GTV_{LN} . (Color figure online)

accordingly. Specifically, a multi-branch network is proposed to adopt a shared encoder and two separate decoders to detect and segment the “tumor-proximal” and “tumor-distal” GTV_{LN} , respectively. A distance-based gating function is designed to generate the corresponding GTV_{LN} sample weights for each branch. By applying the gating function at the outputs of decoders, each branch is specialized to learn the “tumor-proximal” or “tumor-distal” GTV_{LN} features that emulates physician’s diagnosis process. (2) We leverage the early fusion (EF) of three modalities as input to our model, *i.e.*, RTCT, PET and 3D tumor distance map (Fig. 1(bottom row)). RTCT depicts anatomical structures capturing the intensity, appearance and contextual information, while PET provides metastasis functional activities. Meanwhile, the tumor distance map further encodes the critical distance information in the network. Fusion of these three modalities together can effectively boost the GTV_{LN} identification performance. (3) We evaluate on a dataset comprising 651 voxel-wise labeled GTV_{LN} instances in 141 esophageal cancer patients, as the largest GTV_{LN} dataset to date for chest and abdominal radiotherapy. Our method significantly improves the detection mean recall from 72.5% to 78.2%, compared with the previous state-of-the-art lesion detection method [17]. The highest achieved recall of 82.5% is also clinically relevant and valuable. As reported in [5], human observers tend to have relatively low GTV_{LN} sensitivities, *e.g.*, $\sim 80\%$ by even very experienced radiation oncologists. This demonstrates our work’s clinical values.

2 Method

Figure 2 shows the framework of our proposed multi-branch GTV_{LN} detection-by-segmentation method. Similar to [19, 20] which are designed for the pancreatic tumors, we detect GTV_{LN} by segmenting them. We first compute the 3D tumor distance transformation map (Sect. 2.1), based on which any GTV_{LN} is divided into the tumor-proximal or tumor-adjacent subcategory. Next, a multi-branch detection-by-segmentation network is designed where each branch focuses on one subgroup of GTV_{LN} segmentation (Sect. 2.2). This is achieved by applying a binary or soft distance-gating function imposed on the penalty function at the output of the two branches (Sect. 2.3). Hence, each branch can learn specific parameters to specialize on segmenting and detecting the tumor-proximal and tumor-adjacent GTV_{LN}, respectively.

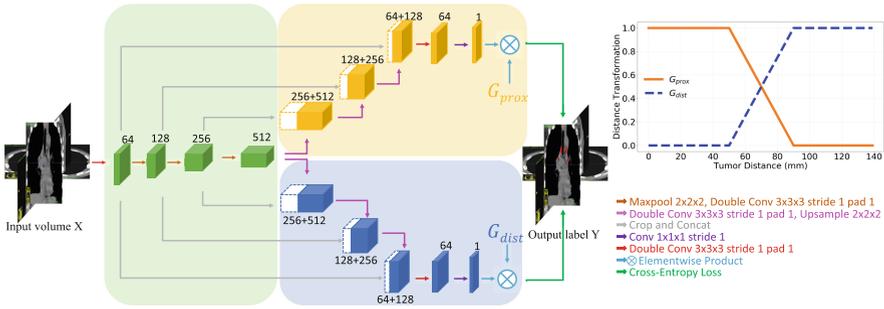


Fig. 2. The overall framework of our proposed multi-branch GTV_{LN} detection and segmentation method. The light green part shows the encoder path, while the light yellow and light blue parts show the two decoders, respectively. The number of channels is denoted either on the top or the bottom of the box. (Color figure online)

2.1 3D Tumor Distance Transformation

To stratify GTV_{LN} into tumor-proximal and tumor-distal subgroups, we first compute the 3D tumor distance transformation map, denoted as \mathbf{X}^D , from the primary tumor \mathcal{O} . The value at each voxel x_i represents the shortest distance between this voxel and the mask of the primary tumor. Let $B(\mathcal{O})$ be a set that includes the boundary voxels of the tumor. The distance transformation value at a voxel x_i is computed as

$$\mathbf{X}^D(x_i) = \begin{cases} \min_{q \in B(\mathcal{O})} d(x_i, q) & \text{if } x_i \notin \mathcal{O} \\ 0 & \text{if } x_i \in \mathcal{O} \end{cases}, \quad (1)$$

where $d(x_i, q)$ is the Euclidean distance from x_i to q . \mathbf{X}^D can be efficiently computed using algorithms such as the one proposed in [9]. Based on \mathbf{X}^D , GTV_{LN} can be divided into tumor-proximal and tumor-distal subgroups using either binary or soft distance-gating function as explained in detail in Sect. 2.3.

2.2 Multi-branch Detection-By-Segmentation via Distance Gating

GTV_{LN} identification is implicitly associated with their distance distributions to the primary tumor in the diagnosis process of physicians. Hence, we divide GTV_{LN} into tumor-proximal and tumor-distal subgroups and conduct detection accordingly. To do this, we design a multi-branch detection-by-segmentation network with each branch focusing on segmenting one GTV_{LN} subgroup. Each branch is implemented by an independent decoder to learn and extract the subgroup specific information, while they share a single encoder to extract the common GTV_{LN} image features. Assuming there are N data samples, we denote a dataset as $\mathbf{S} = \{(\mathbf{X}_n^{\text{CT}}, \mathbf{X}_n^{\text{PET}}, \mathbf{X}_n^{\text{D}}, \mathbf{Y}_n)\}_{n=1}^N$, where \mathbf{X}_n^{CT} , $\mathbf{X}_n^{\text{PET}}$, \mathbf{X}_n^{D} and \mathbf{Y}_n represent the non-contrast RTCT, registered PET, tumor distance transformation map, and ground truth GTV_{LN} segmentation mask, respectively. Without the loss of generality, we drop n for conciseness in the rest of this paper. The total number of branches is denoted as M , where $M = 2$ in our case. A CNN segmentation model is denoted as a mapping function $\mathbb{E} : \mathbf{P} = \mathbf{f}(\mathcal{X}; \Theta)$, where \mathcal{X} is a set of inputs, which consists of a single modality or a concatenation of multiple modalities. Θ indicates model parameters, and \mathbf{P} means the predicted probability volume. Given that $p(y_i|x_i; \Theta_m)$ represents the predicted probability of a voxel $x_i \in \mathcal{X}$ being the labeled class from the m th branch, the overall negative log-likelihood loss aggregated across M branches can be formulated as:

$$\mathcal{L} = \sum_m \mathcal{L}_m(\mathcal{X}; \Theta_m, \mathbf{G}_m) = - \sum_i \sum_m g_{m,i} \log(p(y_i|x_i; \Theta_m)), \quad (2)$$

where $\mathbf{G} = \{\mathbf{G}_m\}_{m=1}^M$ is introduced as a set of volumes containing the transformed gating weights at each voxel based on its distance to the primary tumor. At every voxel $x_i \in \mathbf{G}$, the gating weights satisfies $\sum_m g_{m,i} = 1$.

2.3 Distance-Based Gating Module

Based on the tumor distance map \mathbf{X}^{D} , our gating functions can be designed to generate appropriate GTV_{LN} sample weights for different branches so that each branch specializes on learning the subgroup specific features. In our case, we explore two options: (1) binary distance gating and (2) soft distance gating.

Binary Distance Gating (BG). Based on the tumor distance map \mathbf{X}^{D} , we divide image voxels into two groups, x_{prox} and x_{dis} , to be tumor-proximal and tumor-distal, respectively, where $\text{prox} = \{i|x_i^{\text{D}} \leq d_0, x_i^{\text{D}} \in \mathbf{X}^{\text{D}}\}$ and $\text{dis} = \{i|x_i^{\text{D}} > d_0, x_i^{\text{D}} \in \mathbf{X}^{\text{D}}\}$. Therefore the gating transformations for two decoders are defined as $\mathbf{G}_{\text{prox}} = \mathbf{1}[x_i^{\text{D}} \leq d_0]$ and $\mathbf{G}_{\text{dist}} = 1 - \mathbf{G}_{\text{prox}}$, where $\mathbf{1}[\cdot]$ is an indicator function which equals one if its argument is true and zero otherwise. In this way, we divide the GTV_{LN} strictly into two disjoint categories, and each branch focuses on decoding and learning from one category.

Soft Distance Gating (SG). We further explore a soft gating method that linearly changes the penalty weights of GTV_{LN} samples as they are closer or

further to the tumor. This can avoid a sudden change of weight values when samples are near the proximal and distal category boundaries. Recommended by our physician, we formulate a soft gating module based on \mathbf{X}^D as following:

$$\mathbf{G}_{\text{prox}}(x_i) = \begin{cases} 1 - \frac{x_i^D - d_{\text{prox}}}{d_{\text{dist}} - d_{\text{prox}}} & \text{if } d_{\text{prox}} < x_i^D \leq d_{\text{dist}} \\ 1 & \text{if } x_i^D \leq d_{\text{prox}} \\ 0 & \text{if } x_i^D > d_{\text{dist}} \end{cases}, \quad (3)$$

and $\mathbf{G}_{\text{dist}}(x_i) = 1 - \mathbf{G}_{\text{prox}}(x_i)$ accordingly.

3 Experimental Results

3.1 Dataset and Preprocessing

Dataset. We collected 141 non-contrast RTCTs of esophageal cancer patients, with all undergoing radiotherapy treatments. Radiation oncologists labeled 3D segmentation masks of the primary tumor and all GTV_{LN} . For each patient, we have a non-contrast RTCT and a pair of PET/CT scans. There is a total of 651 GTV_{LN} with voxel-wise annotations in the mediastinum or upper abdomen regions, as the largest annotated GTV_{LN} dataset to-date. We randomly split patients into 60%, 10%, 30% for training, validation and testing, respectively.

Implementation Details. In our experiments, PET scan is registered to RTCT using the similar method described in [6]. Then all coupling pairs of RTCT and registered PET images are resampled to have a consistent spatial resolution of $1 \times 1 \times 2.5$ mm. To generate the 3D training samples, we crop sub-volumes of $96 \times 96 \times 64$ from the RTCT, registered PET and the tumor distance map around each GTV_{LN} as well as randomly from the background. For the distance-gating related parameters, we set $d_0 = 7$ cm as the binary gating threshold, and $d_{\text{prox}} = 5$ cm and $d_{\text{dist}} = 9$ cm as the soft gating thresholds, respectively, as suggested by our clinical collaborator. We further apply random rotations in the x-y plane within 10° to augment the training data.

Detection-by-segmentation models are trained on two NVIDIA Quadra RTX 6000 GPUs with a batch size of 8 for 50 epochs. The RAdam [8] optimizer with a learning rate of 0.0001 is used with a momentum of 0.9 and a weight decay of 0.0005. For inference, 3D sliding windows with a sub-volume of $96 \times 96 \times 64$ and a stride of $64 \times 64 \times 32$ voxels are processed. For each sub-volume, predictions from two decoders are weighted and aggregated according to the gating transformation \mathbf{G}_m to obtain the final GTV_{LN} segmentation results.

Evaluation Metrics. We first describe the hit criteria, *i.e.*, the correct detection, for our detection-by-segmentation method. For an GTV_{LN} prediction, if it overlaps with any ground-truth GTV_{LN} , we treat it as a hit provided that its estimated radius is similar to the radius of the ground-truth GTV_{LN} within the range of [0.5, 1.5]. The performance is assessed using the mean and max recall (mRecall and $\text{Recall}_{\text{max}}$) at a precision range of [0.10, 0.50] with 0.05 interval,

and the mean free response operating characteristic (FROC) at 3, 4, 6, 8 FPs per patient. These operating points were chosen after confirming with our physician.

Comparison Setups. Using the binary and soft distance-based gating function, our multi-branch GTV_{LN} detection-by-segmentation method is denoted as **multi-branch BG** and **multi-branch SG**, respectively. We compare against the following setups: (1) a single 3D UNet [3] trained using RTCT alone or the early fusion (EF) of multi-modalities (denoted as **single-net** method); (2) Two separate UNets trained with the corresponding tumor-proximal and tumor-distal GTV_{LN} samples and results spatially fused together (our preliminary work [21] denoted as **multi-net BG**); and (3) MULAN [17], a state-of-the-art (SOTA) general lesion detection method on DeepLesion [18] that contains more than 10,000 enlarged LNs.

3.2 Quantitative Results and Discussion

Our quantitative results and comparisons are given in Table. 1. Several observations can be drawn on addressing the effectiveness of our proposed methods. **(1)** The multi-modality input, *i.e.*, early fusion (EF) of RTCT, PET and tumor distance map, are of great benefits for detecting the GTV_{LN}. There are drastic performance improvements of absolute 6.7% and 7.2% in mRecall and mFROC when EF is adopted as compared to using RTCT alone. These results validate that input channels of PET functional imaging and 3D tumor distance transform map are valuable for identifying GTV_{LN}. **(2)** The distance-based gating strategies are evidently effective as the options of **multi-net BG**, **multi-branch BG** and **multi-branch SG** consistently increase the performance. For example, the multi-net BG model achieves 74.7% mRecall and 69.5% mFROC, which is a 1.6% and 1.9% improvement against the best single-net model (where no distance-based stratification is used). The performance further boosts with the network models of multi-branch BG and multi-branch SG, to the highest scores of 78.2% mRecall and 72.4% mFROC achieved by the multi-branch SG.

Table 1. Quantitative results of our proposed methods with the comparison to other setups and the previous state-of-the-art.

Methods	CT	EF	mRecall	Recall _{max}	mFROC	FROC@4	FROC@6
Single-net	✓		0.664	0.762	0.604	0.552	0.675
Single-net		✓	0.731	0.820	0.676	0.667	0.713
Multi-net BG [21]		✓	0.747	0.825	0.695	0.668	0.739
Multi-branch BG (Ours)		✓	0.761	0.845	0.679	0.667	0.716
Multi-branch SG (Ours)		✓	0.782	0.843	0.724	0.729	0.738
MULAN [17]	✓		0.711	0.758	0.632	0.632	0.642
MULAN [17]		✓	0.725	0.781	0.708	0.718	0.720

Multi-branch versus Multi-net. Using the distance-based gating strategy, our proposed multi-branch methods perform considerably better than the **multi-net BG** model. Even our second best model **multi-branch BG**, the mean and maximal recalls have been improved by 1.4% (from 74.7% to 76.1%) and 2.0% (from 82.5% to 84.5%) against the **multi-net BG** model. When the multi-branch framework is equipped with the **soft-gating**, marked improvements of absolute 3.5% and 2.9% in both mRecall and mFROC are observed as compared against to the **multi-net BG** model. This validates the effectiveness of our jointly trained multi-branch framework design, and our intuition that gradually changing GTV_{LN} weights for the proximal and distal branches are more natural and effective. As we recall, the multi-net baseline directly trains two separate 3D UNets [3] targeted to segment each GTV_{LN} subgroup. Considering the limited GTV_{LN} training data (a few hundreds of patients), it can be overfitting prone from the split to even smaller patient subgroups.

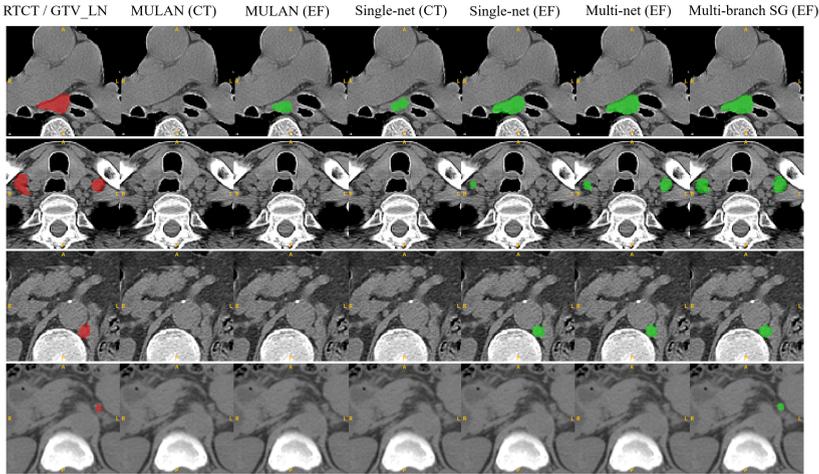


Fig. 3. Four qualitative examples of the detection results using different methods. Red color represents the ground-truth GTV_{LN} overlaid on the RTCT images; Green color indicates the predicted segmentation masks. As shown, for the enlarged GTV_{LN} (top row), most methods can detect it correctly. However, as GTV_{LN} size becomes smaller and contrast is poor, our method can successfully detect them while others struggled. (Color figure online)

Table 1 also compares with the SOTA universal lesion detection method, i.e., MULAN [17] on DeepLesion [16, 18]. We have retrained the MULAN models using both CT and EF inputs, but even the best results, i.e., using EF, have a large gap (72.5% vs. 78.2% mRecall) with our distance-gating networks, which further proves that the tumor distance transformation cue plays a key role in GTV_{LN} identification.

Figure 3 illustrates the visualization results of our method compared to other baselines. For the enlarged GTV_{LN} (top row), most methods can detect it correctly. However, as the size of GTV_{LN} becomes smaller and the contrast is poorer, our method can still successfully detect them while others struggled.

4 Conclusion

In this work, we propose an effective distance-based gating approach in a multi-task deep learning framework to segment GTV_{LN}, emulating the oncologists' high-level diagnosis protocols. GTV_{LN} is divided into two subgroups of "tumor-proximal" and "tumor-distal", by means of binary or soft distance gating. A novel multi-branch detection-by-segmentation network is trained with each branch specializing on learning one subgroup features. We evaluate our method on a dataset of 141 esophageal cancer patients. Our results demonstrate significant performance improvements on the mean recall from 72.5% to 78.2%, as compared to previous state-of-the-art work. The highest achieved GTV_{LN} recall of 82.5% at the 20% precision level is clinically relevant and valuable.

References

1. Barbu, A., Suehling, M., Xu, X., Liu, D., Zhou, S.K., Comaniciu, D.: Automatic detection and segmentation of lymph nodes from CT data. *IEEE Trans. Med. Imag.* **31**(2), 240–250 (2011)
2. Bouget, D., Jørgensen, A., Kiss, G., Leira, H.O., Langø T.: Semantic segmentation and detection of mediastinal lymph nodes and anatomical structures in CT data for lung cancer staging. *Int. J. Comput. Assisted Radiol. surgery*, 14, 1–10 (2019)
3. Çiçek, Ö., Abdulkadir, A., Lienkamp, S.S., Brox, T., Ronneberger, O.: 3D unet: learning dense volumetric segmentation from sparse annotation. In: *MICCAI* (2016)
4. Feulner, J., Zhou, S.K., Hammon, M., Hornegger, J., Comaniciu, D.: Lymph node detection and segmentation in chest CT data using discriminative learning and a spatial prior. *Med. Image Anal.* **17**(2), 254–270 (2013)
5. Goel, R., Moore, W., Sumer, B., Khan, S., Sher, D., Subramaniam, R.M.: Clinical practice in pet/ct for the management of head and neck squamous cell cancer. *Am. J. Roentgenol.* **209**(2), 289–303 (2017)
6. Jin, D., 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). https://doi.org/10.1007/978-3-030-32245-8_21
7. Jin, D., et al.: Deep esophageal clinical target volume delineation using encoded 3D spatial context of tumors, lymph nodes, and organs at risk. In: Shen, D., et al. (eds.) *MICCAI 2019*. LNCS, vol. 11769, pp. 603–612. Springer, Cham (2019). https://doi.org/10.1007/978-3-030-32226-7_67
8. Liu, L., Jiang, H., He, P., Chen, W., Liu, X., Gao, J., Han, J.: On the variance of the adaptive learning rate and beyond. *arXiv preprint arXiv:1908.03265* (2019)
9. Maurer, 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)

10. Network, N.C.C.: NCCN clinical practice guidelines: head and neck cancers. *Am. J. Roentgenol.* Version **2** (2020)
11. Nogues, I., et al.: Automatic lymph node cluster segmentation using holistically-nested neural networks and structured optimization in CT images. In: Ourselin, S., Joskowicz, L., Sabuncu, M.R., Unal, G., Wells, W. (eds.) MICCAI 2016. LNCS, vol. 9901, pp. 388–397. Springer, Cham (2016). https://doi.org/10.1007/978-3-319-46723-8_45
12. Roth, H.R., et al.: Improving computer-aided detection using convolutional neural networks and random view aggregation. *IEEE Trans. Med. Imag.* **35**(5), 1170–1181 (2016)
13. Roth, H.R., et al.: A new 2.5D representation for lymph node detection using random sets of deep convolutional neural network observations. In: Golland, P., Hata, N., Barillot, C., Hornegger, J., Howe, R. (eds.) MICCAI 2014. LNCS, vol. 8673, pp. 520–527. Springer, Cham (2014). https://doi.org/10.1007/978-3-319-10404-1_65
14. Scatarige, J.C., Fishman, E.K., Kuhajda, F.P., Taylor, G.A., Siegelman, S.S.: Low attenuation nodal metastases in testicular carcinoma. *J. Comput. Assisted Tomography* **7**(4), 682–687 (1983)
15. Schwartz, L., et al.: Evaluation of lymph nodes with recist 1.1. *Euro. J. Cancer*, **45**(2), 261–267 (2009)
16. Yan, K., Peng, Y., Sandfort, V., Bagheri, M., Lu, Z., Summers, R.M.: Holistic and comprehensive annotation of clinically significant findings on diverse CT images: learning from radiology reports and label ontology. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pp. 8523–8532 (2019)
17. Yan, K., et al.: MULAN: multitask universal lesion analysis network for joint lesion detection, tagging, and segmentation. In: Shen, D., et al. (eds.) MICCAI 2019. LNCS, vol. 11769, pp. 194–202. Springer, Cham (2019). https://doi.org/10.1007/978-3-030-32226-7_22
18. Yan, K., Wang, X., Lu, L., Summers, R.M.: Deeplesion: automated mining of large-scale lesion annotations and universal lesion detection with deep learning. *J. Med. Imag.* **5**(3), 036501 (2018)
19. Zhu, Z., Lu, Y., Shen, W., Fishman, E.K., Yuille, A.L.: Segmentation for classification of screening pancreatic neuroendocrine tumors. arXiv preprint [arXiv:2004.02021](https://arxiv.org/abs/2004.02021) (2020)
20. Zhu, Z., Xia, Y., Xie, L., Fishman, E.K., Yuille, A.L.: Multi-scale coarse-to-fine segmentation for screening pancreatic ductal adenocarcinoma. In: International Conference on Medical Image Computing and Computer-Assisted Intervention, pp. 3–12. Springer (2019)
21. Zhu, Z., et al.: Detecting scatteredly-distributed, small, and critically important objects in 3d oncology imaging via decision stratification. arXiv preprint [arXiv:2005.13705](https://arxiv.org/abs/2005.13705) (2020)