

LNQ 2023 challenge: Benchmark of weakly-supervised techniques for mediastinal lymph node quantification

Reuben Dorent ¹, Roya Khajavi ¹, Tagwa Idris ², Erik Ziegler ³, Bhanusupriya Somarouthu ^{2,3}, Heather Jacene ^{1,5}, Ann LaCasce ⁵, Jonathan Deissler ¹⁶, Jan Ehrhardt ^{10,11}, Sofija Engelson ¹⁰, Stefan M. Fischer ^{6,7,8}, Yun Gu ¹⁴, Heinz Handels ^{10,11}, Satoshi Kasai ¹², Satoshi Kondo ¹³, Klaus Maier-Hein ^{16,17}, Julia A. Schnabel ^{6,7,8,9}, Guotai Wang ^{15,18}, Litingyu Wang ¹⁵, Tassilo Wald ^{16,17}, Guang-Zhong Yang ¹⁴, Hanxiao Zhang ¹⁴, Minghui Zhang ¹⁴, Steve Pieper ⁴, Gordon Harris ^{2,3}, Ron Kikinis ¹, Tina Kapur ¹

1 Brigham and Women's Hospital, Harvard Medical School, Boston, MA, USA; **2** Massachusetts General Hospital, Harvard Medical School, Boston, MA, USA; **3** Yunu, Inc., Cary, NC, USA; **4** Isomics Inc, Cambridge, MA, USA; **5** Dana-Farber Cancer Institute, Boston, MA, USA; **6** Technical University Munich, Munich, Germany; **7** Institute of Machine Learning in Biomedical Imaging, Helmholtz Munich, Munich, Germany; **8** Munich Center of Machine Learning (MCML), Munich, Germany; **9** School of Biomedical Engineering and Imaging Sciences, King's College London, London, UK; **10** Institute of Medical Informatics, University of Lübeck, Lübeck, Germany; **11** German Research Center for Artificial Intelligence, Lübeck, Germany; **12** Niigata University of Health and Welfare, Niigata, Japan; **13** Muroran Institute of Technology, Hokkaido, Japan; **14** Institute of Medical Robotics, Shanghai Jiao Tong University, Shanghai, China; **15** University of Electronic Science and Technology of China, Chengdu, China; **16** Division of Medical Image Computing, German Cancer Research Center (DKFZ), Heidelberg, Germany; **17** University of Heidelberg, Heidelberg, Germany; **18** Shanghai AI laboratory, Shanghai, China.

Abstract

Accurate assessment of lymph node size in 3D CT scans is crucial for cancer staging, therapeutic management, and monitoring treatment response. Existing state-of-the-art segmentation frameworks in medical imaging often rely on fully annotated datasets. However, for lymph node segmentation, these datasets are typically small due to the extensive time and expertise required to annotate the numerous lymph nodes in 3D CT scans. Weakly-supervised learning, which leverages incomplete or noisy annotations, has recently gained interest in the medical imaging community as a potential solution. Despite the variety of weakly-supervised techniques proposed, most have been validated only on private datasets or small publicly available datasets. To address this limitation, the Mediastinal Lymph Node Quantification (LNQ) challenge was organized in conjunction with the 26th International Conference on Medical Image Computing and Computer Assisted Intervention (MICCAI 2023). This challenge aimed to advance weakly-supervised segmentation methods by providing a new, partially annotated dataset and a robust evaluation framework. A total of 16 teams from 5 countries submitted predictions to the validation leaderboard, and 6 teams from 3 countries participated in the evaluation phase. The results highlighted both the potential and the current limitations of weakly-supervised approaches. On one hand, weakly-supervised approaches obtained relatively good performance with a median Dice score of 61.0%. On the other hand, top-ranked teams, with a median Dice score exceeding 70%, boosted their performance by leveraging smaller but fully annotated datasets to combine weak supervision and full supervision. This highlights both the promise of weakly-supervised methods and the ongoing need for high-quality, fully annotated data to achieve higher segmentation performance.

Keywords

Machine Learning, Image Segmentation, Weak Supervision, Lymph Node

Article informations

<https://doi.org/10.59275/j.melba.2025-d482>

Received: 2024-08-20, Published 2025-01-29

Corresponding author: reuben.dorent@inria.fr

Special issue: **MICCAI 2023 LNQ challenge special issue**

Guest editors: Steve Pieper, Erik Ziegler, Tawa Idris, Bhanusupriya Somarouthu, Reuben Dorent, Gordon Harris, Ron Kikinis

©2025 Dorent, et al.. License: CC-BY 4.0



1. Introduction

Machine learning has recently achieved outstanding medical image segmentation performance. While initial frameworks were tailored to specific segmentation tasks, significant progress has been made in developing network architectures and training procedures that are robust and applicable across segmentation tasks (Isensee et al., 2021; Cardoso et al., 2022). These robust frameworks generally rely on the availability of large, fully annotated datasets during training. This reliance, however, is often impractical. Manual medical data segmentation is a labor-intensive and time-consuming process that demands rare expertise, leading to small, fully annotated datasets. Moreover, obtaining complete annotations in an image is sometimes unfeasible due to factors such as the disproportionate size of the targeted object of interest compared to the image (e.g., cells in histopathology images), data ambiguity (e.g., tumor in ultrasound imaging), or the presence of a very large number of objects to be segmented (e.g., lymph nodes in 3D scans).

In response to these challenges, weakly-supervised learning methods have emerged as promising alternatives. These approaches leverage sparse, noisy, or incomplete annotations to train machine learning models, significantly reducing the need for complete annotations at training time. In the context of medical imaging, weakly-supervised learning can utilize various forms of weak supervision, such as image-level labels (Ouyang et al., 2019), bounding boxes (Kervadec et al., 2020), scribbles (Zhang and Zhuang, 2022; Dorent et al., 2020), points (Roth, Holger and Zhang, Ling and Yang, Dong and Milletari, Fausto and Xu, Ziyue and Wang, Xiaosong and Xu, Daguang, 2019; Dorent et al., 2021b; Can et al., 2018), linear measurements (Cai et al., 2018; Li and Xia, 2020) or partial annotations (Mehrtash et al., 2024; Dorent et al., 2021a), to achieve competitive performance in segmentation tasks.

To benchmark new and existing weakly-supervised techniques for medical image segmentation, we organized the Mediastinal Lymph Node Quantification (LNQ) challenge in conjunction with the 26th International Conference on Medical Image Computing and Computer Assisted Intervention (MICCAI 2023). The challenge’s goal was to segment and identify mediastinal lymph nodes in contrast-enhanced computed tomography (CT) scans using a new large and partially annotated dataset. Note that participants were allowed to exploit existing smaller, publicly available, and fully annotated datasets.

Accurate lymph node size estimation is critical for staging cancer patients, initial therapeutic management, and, in longitudinal scans, assessing response to therapy. Current standard practice for quantifying lymph node size is based on various criteria that use unidirectional or bidirectional

measurements on just one or a few nodes, typically on just one axial slice. These evaluations are performed on routine CT scans. However, humans have hundreds of lymph nodes, any number of which may be enlarged to various degrees due to disease or immune response. While a normal lymph node may be approximately 5 mm in diameter, a diseased one may be several cm in diameter. The mediastinum, the anatomical area between the lungs and around the heart, may contain ten or more lymph nodes, often with three or more enlarged greater than 1 cm in afflicted patients. Accurate volumetric assessment would thus provide information to evaluate lymph node disease and provide better sensitivity to detect volumetric changes indicating response to therapy.

While automated full 3D segmentation of all abnormal lymph nodes could improve cancer treatment, only small, fully annotated datasets are currently publicly available to train machine learning frameworks. For these reasons, we proposed a weakly-supervised benchmark that aims to automatically perform lymph node segmentation in 3D CT scans using partial annotations. Specifically, we partially annotated (a few nodes only) a large number of scans and evaluated the performance of the participants’ methods using a large, fully annotated dataset.

This paper summarizes the LNQ 2023 challenge and is structured as follows. First, a review of existing datasets used to perform lymph node quantification is presented in Section 2. Then, the design of the LNQ challenge is given in Section 3. Section 4 presents the evaluation strategy of the challenge (metrics and ranking scheme). Participating methods are then described and compared in Section 5. Finally, Section 6 presents the results of the participating teams, and Section 7 provides a discussion and concludes the paper.

2. Related Works

We performed a literature review to survey the datasets used to assess lymph node quantification. Many lymph node quantification techniques have been explored and validated on private datasets, for example, Tan et al. (2018); Barbu et al. (2012); Stapleford et al. (2010); Feulner et al. (2013). Since these datasets are private and working implementations of the methods are not available, it is impossible to benchmark existing and new methods with these techniques.

Other authors have used public datasets to validate their methods. We present these open datasets and highlight their limitations for evaluating lymph node quantification:

- *CT Lymph Node dataset*: The NIH CT Lymph Node dataset (Roth et al., 2015) comprises a total of 176 contrast-enhanced CT series from 176 patients. Among

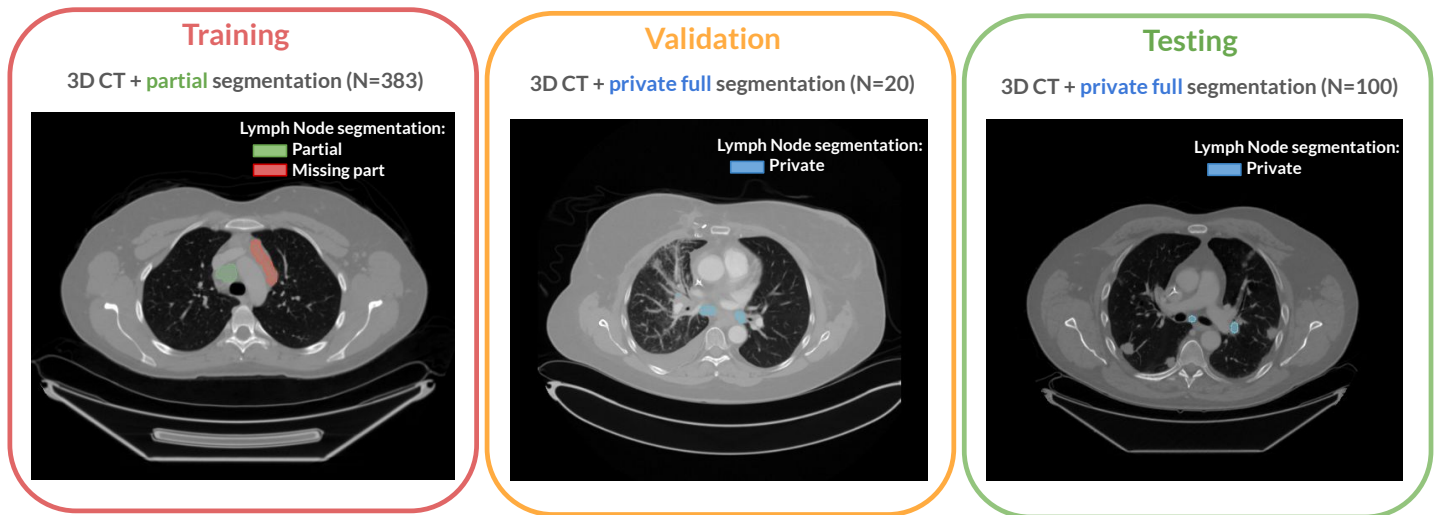


Figure 1: Overview of the challenge dataset. Only partial annotations (green) were made available for the training 3D CT scans. Missing training data nodes are shown in red. Full segmentations (blue) of all nodes were performed for evaluation on the validation and testing sets, and remained private.

them, 89 CT volumes were obtained at the chest level (mediastinum). Partial node segmentation is provided, corresponding to 387 nodes with a short axis diameter (SAD) ≥ 1 cm, which is considered clinically enlarged and abnormal. Bouget et al. (2023) extended and refined these annotations for all the mediastinal lymph nodes in these 89 volumes. All suspicious regions are segmented as lymph nodes, including nodes with short-axis measurement less than 1 cm. In total, 2912 nodes were segmented.

- *St. Olavs Hospital dataset*: This dataset (Bouget et al., 2019) comprises 15 contrast-enhanced CT volumes from 15 patients with confirmed lung cancer diagnoses. Segmentation of all the lymph nodes is provided. A total of 384 lymph nodes were annotated in this dataset.

Moreover, other datasets that comprise contrast-enhanced CT scans have been released from patients with lung cancer, such as NSCLC-Radiomics (Aerts et al., 2014), NSCLC Radiogenomics (Bakr et al., 2018), NSCLC-Radiomics Interobserver1 (Wee et al., 2019), RIDER Lung CT (Zhao et al., 2009). However, these datasets do not provide any segmentations of lymph nodes.

In conclusion, test sets used to assess segmentation methods for mediastinal lymph node segmentation are either private, relatively small, or partially annotated.

On the methodological side, a wide range of weakly-supervised methods have been proposed. Nonetheless, a thorough review of weakly-supervised methodologies is out of the scope of this paper, except to note that it is an active area of research with many unsolved challenges. We refer the interested reader to Zhang et al. (2020) for a review of these.

3. Challenge description

3.1 Overview

The goal of the LNQ challenge was to benchmark new and existing weakly-supervised techniques for lymph node quantification. The proposed segmentation task focused on segmenting all the lymph nodes in contrast-enhanced computed tomography (CT) scans. Participants had access to a training set of partially annotated CT scans. Participant's algorithms were evaluated on a fully annotated dataset.

3.2 Data description

3.2.1 Data overview

The challenge cohort is a cross-institutional dataset of 513 chest CT scans acquired during patient treatment for various cancer types. The dataset originates from the Tumor Imaging Metrics Core (TIMC), a multi-institutional imaging core lab (Massachusetts General Hospital, Dana-Farber Cancer Institute, Brigham and Women's Hospital, Boston, MA, USA) that provides multimodality imaging measurements to evaluate treatment response in patients enrolled in oncology clinical trials. This dataset was used for training, validation, and testing. Figure 1 presents an overview of the challenge dataset.

The complete LNQ dataset (training, validation, and testing) contained CT images collected on 513 patients (Male:Female 239:274) enrolled in oncology clinical trials from 2007-2020. For each patient, contrast-enhanced CT scans were acquired. The Institutional Review Board at the Mass General Brigham (MGB IRB) approved the protocol (2020P000211), including public sharing of the data

Table 1: Summary of characteristics of the LNQ data sets. Median [Q1-Q3] are given for numerical values.

| | Training | Validation | Test |
|---------------------|------------------------------------|------------------------------------|------------------------------------|
| Modality | contrast-enhanced CT | contrast-enhanced CT | contrast-enhanced CT |
| Number of scans | 383 | 20 | 100 |
| Number of patients | 383 | 20 | 100 |
| Cancer type | Breast (18%) | Breast (25%) | Breast (4%) |
| | Chronic lymphocytic leukemia (15%) | Chronic lymphocytic leukemia (30%) | Chronic lymphocytic leukemia (11%) |
| | Hodgkin lymphoma (7%) | Hodgkin lymphoma (5%) | Hodgkin lymphoma (8%) |
| | Lung non-small cell (9%) | Lung non-small cell (25%) | Lung non-small cell (16%) |
| | Lung small cell (6%) | Lung small cell (5%) | Lung small cell (7%) |
| | Renal cell (5%) | Renal cell (0%) | Renal cell (7%) |
| | Others (40%) | Others (10%) | Others (47%) |
| Public annotations | Some lymph nodes | × | × |
| Private annotations | × | All lymph nodes | All lymph nodes |
| In-plane matrix | 512 (100%) | 512 (100%) | 512 (100%) |
| Slice number | 114 [96-136] | 112 [92-127] | 86 [77-93] |
| In-plane res. in mm | 0.8 [0.8-0.9] | 0.8 [0.8-0.9] | 0.9 [0.8-0.9] |
| Slice spacing in mm | 3.0 [2.5-3.8] | 3.0 [2.5-3.8] | 3.8 [3.8 - 5.0] |
| Male:Female | 47% : 53% | 45% : 55% | 47% : 53% |

(2020P003754).

The dataset comprises patients with various types of cancer. The top six primary patient cancers, which accounted for 60% of all patients, were breast cancer (N=80), chronic lymphocytic leukemia (CLL) (N=74), non-small cell lung cancer (N=58), Hodgkin’s lymphoma (N=38), small cell lung cancer (N=30), and renal cell cancer (N=25). The remaining patient cancer types, accounting for 40% of all patients, including thyroid cancer, adenocarcinoma, endometrial adenocarcinoma, melanoma, head and neck cancer, non-Hodgkin’s lymphoma, prostate cancer, mesothelioma, esophageal cancer, ovarian cancer, and colon cancer.

The radiology and oncology experts on our team have determined that for this challenge, the CT appearance of clinically important lymph nodes is not influenced by the primary cancer type. Based on this, we have included as many cases as possible with segmented lymph nodes of any primary cancer type on the premise that access to more data that experts judge to be similar in appearance is more likely to provide a robust segmentation model.

The dataset was randomly split into training (76%), validation (4%) and test (20%) sets. Table 1 shows the distribution of features (patient sex, cancer, slice thickness, and in-plane resolution) across these sets.

3.2.2 Image acquisition

Images from the LNQ challenge were acquired with routine clinical CT scanners from various manufacturers (e.g. GE Healthcare Discovery CT750HD, GE Medical System BrightSpeed, Siemens SOMATOM Definition, Toshiba Aquilion, Philips iCT). The scans are of routine clinical quality, with 2 to 5mm slice spacing and an in-plane resolution of approximately 1×1 mm or smaller. Acquisitions

are typically 512x512 axial scans with 100 or more slices (range: 48-656 slices).

3.2.3 Annotation protocol

All imaging datasets were manually segmented following two annotation protocols. While the training set was partially annotated, the validation and test sets were fully annotated, following the considerations described below.

The training set was annotated based on data from the clinical trials workflow. This annotation process involves two steps: lymph node selection and manual volumetric segmentation. Since the initial annotations are created within the context of clinical trials, the cases are not initially fully annotated; only specific lesions or lymph nodes are marked. Specifically, expert image analysts used the Yunu (Cary, NC, previously known as Precision Imaging Metrics (PIM)) clinical trials imaging informatics software system to select lesions, focusing on those to be followed with bi-dimensional measurements in single axial slices of longitudinal CT scans. Nodes were selected according to trial protocols defined by the sponsor, which may vary for different patients. This means the selected lesions may not always be the largest present in the case (Eisenhauer et al., 2009). Moreover, protocols may have different selection criteria or only call for annotations of a maximum number of lesions per patient. In the second step, the volumetric extent of these selected nodes was manually segmented to create the weak annotations for the heterogeneous training CT data of the challenge.

The validation and test sets were fully segmented to reflect the full mediastinal lymph node disease burden and used to assess the performance of the LNQ participants’ methods. All lymph nodes considered abnormal, with an estimated short axis length larger than 1 cm, were tar-

geted for segmentation. Surrounding structures excluded from the lesion boundary included large vessels, artifacts, and non-nodal components. We acknowledge that visual assessment of the short axis is error-prone, so this criterion may not always be satisfied for nodes with a short axis around the threshold.

Volumetric segmentations were performed using 3D Slicer (Pieper et al., 2004; Fedorov et al., 2012). All lymph nodes were segmented by a trained radiologist (TI), with the support of an instructor in radiology (BS), and in consultation with a senior imaging specialist (HJ). For each patient, the CT scan and the bidimensional measurements of the target lymph node were loaded into 3D Slicer (version 4.11). The Segment Editor module was then used to manually delineate the lymph node boundary at the native resolution of the CT scan. The Draw tool within the Editor module was employed to draw freehand boundaries on axial cross-sections while referencing the sagittal and coronal planes.

3.2.4 Data curation

Data were fully de-identified by removing all personal health information identifiers and creating fresh DICOM files containing only approved tags and images. The MGB IRB (2020P000211 and 2020P003754), approved the de-identification procedure, and all header data was manually reviewed by MGB staff. Image series were manually reviewed for quality and to ensure no issues were caused by de-identification.

Images and segmentation masks were distributed as NRRD files (.nrrd). The training and validation datasets were made publicly available via the challenge page ¹. In contrast, the testing set test set remained private. As we expect this annotated test dataset to be used for other purposes, the data has been on the Cancer Imaging Archive (TCIA) (Idris et al., 2024).

3.3 Challenge setup

The validation phase took place on Grand Challenge, a renowned platform for biomedical challenges, which facilitated automated validation leaderboard management. Participant submissions were automatically assessed using the `evalutils` and `MedPy` Python packages. Each participant could make up to three daily submissions on the validation leaderboard. This phase ran from May 1, 2023, to August 30, 2023.

Following best practice guidelines for organizing challenges, the test set remained private to help ensure fairness of the evaluation. Participants were required to containerize their algorithms using Docker, in accordance with Grand

Challenge guidelines, and submit their Docker containers for evaluation on the test set. Only one submission was permitted for the test set evaluation. Participants were first encouraged to test their Docker containers on the validation set without submission limits to ensure the algorithms were containerized correctly. If the predictions matched those generated on their machines, participants could submit their algorithms for test set evaluation.

4. Metrics and evaluation

The choice of metrics to evaluate participants' algorithms and the ranking strategy are crucial for accurate interpretation and reproducibility of results (Maier-Hein et al., 2018). In this section, we adhere to the BIAS best practice recommendations for evaluating challenges (Maier-Hein et al., 2020).

4.1 Choice of the metrics

The primary characteristic to optimize for the algorithms is prediction accuracy. Since relying on a single metric for segmentation assessment can result in less robust rankings, we selected two metrics: the Dice similarity coefficient (DSC) and the Average symmetric surface distance (ASSD). These metrics are widely used in previous challenges (Kavur et al., 2021; Antonelli et al., 2022; Dorent et al., 2023) due to their simplicity, rank stability, and effectiveness in evaluating segmentation accuracy.

Let S represent the predicted binary segmentation mask of the lymph nodes, and G represent the manual segmentation. The Dice Score coefficient measures the similarity between masks S and G by normalizing the size of their intersection over the average size of the masks:

$$\text{DSC}(S, G) = \frac{2 \sum_i S_i G_i}{\sum_i S_i + \sum_i G_i} \quad (1)$$

Let B_S and B_G be the boundaries of the segmentation mask S and the manual segmentation G . The average symmetric surface distance (ASSD) is calculated as the average of all Euclidean distances (in mm) from points on boundary B_S to the boundary B_G and vice versa:

$$\text{ASSD}(S, G) = \frac{\sum_{s_i \in B_S} d(s_i, B_G) + \sum_{s_i \in B_G} d(s_i, B_S)}{|B_S| + |B_G|} \quad (2)$$

where d denotes the Euclidean distance.

4.2 Ranking scheme

We employed a standard ranking scheme, successfully used in other challenges such as BraTS (Bakas et al., 2018) and crossMoDA (Dorent et al., 2023). Teams are ranked for each test case and each metric (DSC and ASSD). In the

1. <https://lnq2023.grand-challenge.org/data/>

case of ties, the lowest rank is assigned to the tied values. The overall rank score is calculated by first averaging individual rankings across all cases (cumulative rank) and then averaging these cumulative ranks across all patients for each team. The final team rankings are based on these rank scores. This ranking scheme was defined and published before the challenge began and was available on the Grand Challenge page².

To assess the stability of the ranking scheme, we used the bootstrapping method described by Wiesenfarth et al. (2021). We generated 1,000 bootstrap samples by randomly drawing 100 test cases with replacements from the test set, where each sample retained approximately 63% of distinct cases. The ranking scheme was then applied to each bootstrap sample. We compared the original test set ranking to the rankings from individual bootstrap samples using Kendall's τ , which ranges from -1 (reverse order) to 1 (identical order).

5. Participating methods

A total of 208 teams registered for the challenge, allowing them to download the data. 16 teams from 5 different countries submitted predictions to the validation leaderboard. Among them, 6 teams from 3 countries submitted their containerized algorithm for the evaluation phase.

In this section, we summarize the methods used by these 6 teams. Each method is assigned a unique color code used in the tables and figures. Brief comparisons of the proposed techniques in terms of methodology and implementation details are presented in Table 2.

● **Skeleton Suns (1st place, Deissler et al.)** The authors implemented their method using the nnU-Net architecture (Isensee et al., 2021) with the extension of a residual encoder (Isensee and Maier-Hein, 2019). The overall learning strategy involved a multi-step approach to address the challenges of partially annotated datasets. Initially, the nnU-Net framework was trained using the LNQ challenge dataset (Khajavi et al., 2023). Due to incomplete annotations in the training data, strategic data enhancement was employed, and an additional fully labeled dataset from TCIA (Roth et al., 2015) with refined segmentations (Bouget et al., 2023) was incorporated to improve training efficacy. To handle incomplete annotations, the team used TotalSegmentator (Wasserthal et al., 2023) to identify and label non-lymph node structures, effectively refining the segmentation labels. The remaining unlabeled regions were excluded from the loss calculation to enable the model to predict unlabeled lymph nodes within these. Additionally, the Bodypartregression toolkit (Schuhegger, 2021) focused the model on the mediastinal region, assigning anatomical

scores to each axial slice to exclude non-target areas. These preprocessing steps ensured the model trained effectively on confirmed labels without mislabeling background areas as lymph nodes. The network employed was the '3d fullres' configuration of nnU-Net, with adaptations including adjustments to batch size, patch size, and an extended data augmentation strategy. The training involved a 5-fold cross-validation process, with the final model being an ensemble of these cross-validated models. The ensemble predictions were averaged during inference, though test time augmentations were omitted due to time constraints, slightly degrading performance. Post-processing steps included averaging softmax outputs from the ensemble models to produce the final segmentation masks.

● **IMR (2nd place, Zhang et al.)** The team also proposed a semi- and weakly-supervised learning method for automatically segmenting clinically relevant lymph nodes in the mediastinal area of contrast-enhanced CT scans that utilized both partial annotation and full annotation data for two-stage training. First, to better capture the anatomic and semantic representations of mediastinal lymph nodes, a pre-processing approach guided by lung masks and airway maps was used to crop mediastinal VOIs, which included four steps: 1) the lung mask was extracted to obtain the initial VOI boundary of the lung (Hofmanninger et al., 2020); 2) this initial VOI was used to segment the airway map (Zhang et al., 2023); 3) airway voxels were removed in the lung mask to extract the secondary VOI boundary; and 4) the final VOI input was based on the two VOI bounding boxes with margin settings. Then, a two-stage pipeline was designed for semantic segmentation that benefits from partial annotation data. In the first stage, a full-resolution nnU-Net (Isensee et al., 2021) model was trained initially from scratch with full annotation data from the CT Lymph Node dataset (Roth et al., 2015) with refined annotations (Bouget et al., 2023) and the St. Olavs Hospital dataset (Bouget et al., 2019) based on one-fold of the 5-fold cross-validation for 1000 epochs. In the second stage, the trained model predicted pseudo labels for LNQ training data, which were combined with their partial annotations to produce new lymph node labels. The final model was finetuned using jointly full annotation data and updated partial annotation data for 300 epochs. To exclude non-diseased lymph nodes, each individual component whose volume was less than the volume of a sphere with a radius of 5 mm was removed in the post-processing step.

● **CompAI (3rd place, Fischer et al.)** This team employed a semi- and weakly-supervised approach. The authors proposed incorporating the TotalSegmentator toolbox (Wasserthal et al., 2023) to generate pseudo labels and loss masking for handling incomplete annotations in

2. <https://lnq2023.grand-challenge.org/>

Table 2: Comparisons of the proposed techniques in terms of methodology and implementation details.

| | ● Skeleton Suns | ● IMR | ● CompAI | ● Hilab | ● IMI | ● SKJP |
|-----------------------------|--|--|--|-----------------------------------|--|-----------------------------|
| Network architecture | 3D U-Net (nnU-Net) | 3D U-Net (nnU-Net) | 3D U-Net (nnU-Net) | 3D U-Net (nnU-Net) | 3D V-Net (nnU-Net) | 2.5D U-Net |
| Use of external datasets | TCIA CT Lymph Nodes + Bouget refinements | TCIA CT Lymph Nodes + Bouget refinements St. Olavs Hospital | TCIA CT Lymph Nodes + Bouget refinements NSCLC datasets | × | TCIA CT Lymph Nodes + Bouget refinements St. Olavs Hospital | × |
| Cropping | × | Lung and airway | Lung | Lung | Lung | × |
| Weakly Supervision Approach | Background pseudo-labeling | Self-supervision | Background pseudo-labeling | Partial-supervised loss functions | Probabilistic atlas | × |
| Post-processing | × | Small component removal | Small component removal | Small component removal | Small component removal | Largest component selection |

combination with supervised learning using the nnU-Net framework (Isensee et al., 2021). Alongside the challenge training data, the team utilized the fully annotated TCIA CT Lymph Nodes dataset (Roth et al., 2015), providing annotations of all pathologic lymph nodes. Instead of training solely on pathologic lymph nodes, they replaced original annotations with refined annotations containing all visible lymph nodes (Bouget et al., 2023). Furthermore, to increase the dataset size, the authors included the public lung cancer datasets NSCLC radiomics (Aerts et al., 2014), NSCLC radiogenomics (Bakr et al., 2018), and NSCLC interobserver (Wee et al., 2019). First, they created ROIs by cropping each volume to the lung bounding box via the TotalSegmentator. To handle incomplete annotations, pseudo labels were generated from TotalSegmentator structures. Those structures should, by definition, not contain any lymph nodes. Labeling those structures as background increased the training supervision significantly. The authors also masked the loss of remaining unlabeled voxels from the training process. Only one nnU-Net instance was trained on the preprocessed data with adjusted hyperparameters for learning rate and intensity clipping. Furthermore, the authors applied a postprocessing step on the segmentation output by discarding all lymph node components with a shortest axis diameter less than 10 mm. In their final challenge report (Fischer et al., 2024), the authors show that this postprocessing hurt the overall performance and was based on their misinterpreting the challenge goal.

● **Hilab (4th place, Wang et al.)** This team proposed a framework that combines the techniques of self- and weakly-supervised learning. First, they utilized a self-supervised method named Model Genesis (Zhou et al., 2021) to initialize the weights of a VNet-like (Milletari et al., 2016) model with one encoder and two decoders. Second, they proposed to add perturbation to the input of one of the decoders to generate variability between the predictions from each decoder. These two predictions were then dynamically mixed to produce better pseudo labels. To better utilize incomplete annotations, Partial Cross-Entropy (PCE) loss (Lee and Jeong, 2020) and Tversky loss (Salehi et al., 2017) were applied to this task to bal-

ance the supervisory signals of the foreground and background. They also used noise-robust Symmetric Cross-Entropy (SCE) loss (Wang et al., 2019) to further extract information from incomplete annotations. To prevent the erroneous foreground voxels in the pseudo labels from misleading the model training, a weighted Cross-Entropy loss was employed following the approach outlined in Zheng and Yang (2021). Only the LNQ challenge training data were used and cropped to the lung region during inference and training. The model was trained in a patch-based manner and inferred using a sliding window strategy. Only the prediction from the decoder without perturbation was used as the inference result. Then, the predicted foreground was refined based on the actual volume of the connected domain, the intensity of the pixel value, and whether it is at the edge of the image. Finally, the prediction was re-sampled to its original size. The authors later found in their challenge report (Wang et al., 2024) that although the SCE loss and self-learning method improved the performance on the validation set, degradation on the test set was observed.

● **IMI(5th place, Engelson et al.)** The team proposed an ensemble of five segmentation models based on full-resolution nnU-Net trained on multiple anatomical priors as additional input (Isensee et al., 2021; Engelson et al., 2024b). To address the challenges arising from the weak annotations of the training data, a probabilistic lymph node atlas was registered to the training data to identify regions with high lymph node occurrence, which was then used for loss weighting and post-processing. Further, the authors addressed the heterogeneity of the training data and lymph node appearance by using a strong augmentation called GIN & IPA introduced by Ouyang et al. (2021). Preprocessing contained the generation of the anatomical priors using atlas-to-patient registration based on selected segmentation masks from the TotalSegmentator algorithm (Wasserthal et al., 2023), cropping to the lung region, and normalization. For post-processing, the threshold for binarization was lowered according to the lymph node atlas, and segmentation masks below a minimum diameter size of 5 mm were removed. The input data consisted of the weakly-labeled LNQ 2023 training data and publicly available CT

Lymph Node dataset (Roth et al., 2015) with refined annotations Bouget et al. (2023) and the St. Olavs Hospital dataset (Bouget et al., 2019). The fully annotated data was oversampled during training. After the publication of the fully annotated LNQ 2023 validation and test data, the authors realized that the restriction to a minimum lymph node size degraded segmentation accuracy and that refining on the fully annotated datasets improved performance, as mentioned in their report (Engelson et al., 2024a).

- SKJP (6th place, Kondo, et al.)** The team used encoder-decoder type deep neural networks. Although the input images are 3D volumes, they use a 2D deep neural network model with multi-slice inputs (2.5D CNNs). Specifically, they used a 2D U-Net as the segmentation network and replaced its encoder part with EfficientNet (Tan and Le, 2019). In their model, N consecutive slices in each input volume were concatenated and treated as an N channel 2D input image. The input image was then processed to produce one slice in a segmentation mask volume. The input slices were along the transverse plane. During training, slices were randomly selected M times from each volume in the training dataset in each epoch. The loss function was a sum of Dice loss and cross-entropy loss with equal weights. The dataset was split into 353 training samples and 37 validation samples. The optimizer was AdamW and the learning rate was decreased at every epoch with cosine annealing. The model was trained for 100 epochs, and the version with the lowest validation loss was selected as the final model. Random intensity shifts and random affine transformations were applied as data augmentation. Hyperparameter tuning was conducted, including adjustments to the encoder size, number of slices, and initial learning rate. As a result, EfficientNet-B7 was chosen as the encoder, with 5 slices and an initial learning rate of 10^{-3} . During inference, each volume was processed slice by slice, and post-processing involved selecting the largest connected component.

6. Results

Participants were required to submit their algorithm by 20th September 2023. The final results on the testing set were announced during the LNQ workshop at the MICCAI 2023 conference. This section presents the results obtained by the participant teams on the test set and analyses the stability and robustness of the proposed ranking scheme.

6.1 Overall segmentation performance

The final scores for the 6 teams are reported in Table 3 in the order in which they ranked. Figures 2a and 2b show the box plots for each metric (Dice and ASSD) and are color-coded according to the team.

The winner of the LNQ2023 challenge was the Skeleton Suns, with a rank score of 2.7. Skeleton Suns is the only team with a median ASSD lower than 3 mm. Other teams in the top three also obtained encouraging results with a median DSC greater than 69%. In contrast, the low DSC and high ASSD scores of the team with the lowest rank highlight the complexity of the task.

The top three teams employed a semi- and weakly-supervised approach, combining full supervision using existing fully annotated datasets and weak supervision using the LNQ dataset to train their frameworks. In contrast, methods that only leverage the LNQ weak labels (●●) underperformed compared to these semi-supervised approaches. As shown in Table 3, their medians are significantly lower, and their interquartile ranges (IQRs) are larger. This highlights the effectiveness of leveraging complete annotations in combination with partial annotations to improve model performance.

6.2 Analysis of the variations in performance

While the performance achieved by the top-performing team is high on average, there are patients for whom the best models do not perform well. In this section, we aim to analyze the robustness of the proposed models and identify the more challenging cases.

First, we propose determining whether low-performing cases are similarly distributed across all teams. Specifically, we compare the performance of the second and third teams with that of the top-performing team for each test case. The Figure 3 shows the relationship between the Dice scores of the different techniques, with the x -axis representing the Dice score of the best-performing team and the y -axis representing the Dice scores of the other teams. It is clear that in general challenging cases for the top-performing team are also difficult for the other teams. Additionally, the first two teams exhibit very similar performance for each case, while the third team struggles with some cases successfully segmented by the other two. This suggests that some cases were harder to segment.

To determine if certain patient attributes are predictive of segmentation difficulty, we analyzed the distribution of the best Dice scores across all methods for each case, considering key patient attributes such as cancer type and sex. The results are shown in Figure 4. We found that the best Dice scores had a similar distribution across sexes, as depicted in Figure 4a. Conversely, Figure 4b shows the distribution of scores by patient cancer type. It can be observed that patients with Hodgkin lymphoma had statistically lower scores compared to those with chronic lymphocytic leukemia (CLL) ($p = 0.009$), renal cell carcinoma ($p = 0.003$), breast cancer ($p = 0.008$), and other types ($p = 0.27$). This could be attributed to the fact that

Table 3: Metrics values and corresponding scores of submission. Median and interquartile values are presented. The best results are given in bold. Arrows indicate the favorable direction of each metric.

| | Challenge Rank | | Lymph Node | |
|-----------------|----------------|--------------|---------------------------|---------------------------|
| | Global Rank ↓ | Rank Score ↓ | DSC (%) ↑ | ASSD (mm) ↓ |
| ● Skeleton Suns | 1 | 2.1 | 71.2 [63.4 - 78.6] | 2.95 [2.13 - 5.05] |
| ● IMR | 2 | 2.6 | 70.0 [63.7 - 77.1] | 4.15 [2.92 - 6.3] |
| ● CompAI | 3 | 2.9 | 69.0 [57.1 - 74.8] | 5.05 [3.12 - 8.79] |
| ● Hilab | 4 | 3.9 | 61.0 [54.0 - 71.1] | 5.92 [3.77 - 8.81] |
| ● IMI | 5 | 3.9 | 61.2 [46.3 - 71.1] | 6.0 [3.84 - 9.43] |
| ● SKJP | 6 | 5.5 | 43.3 [22.6 - 56.6] | 10.4 [6.56 - 15.39] |

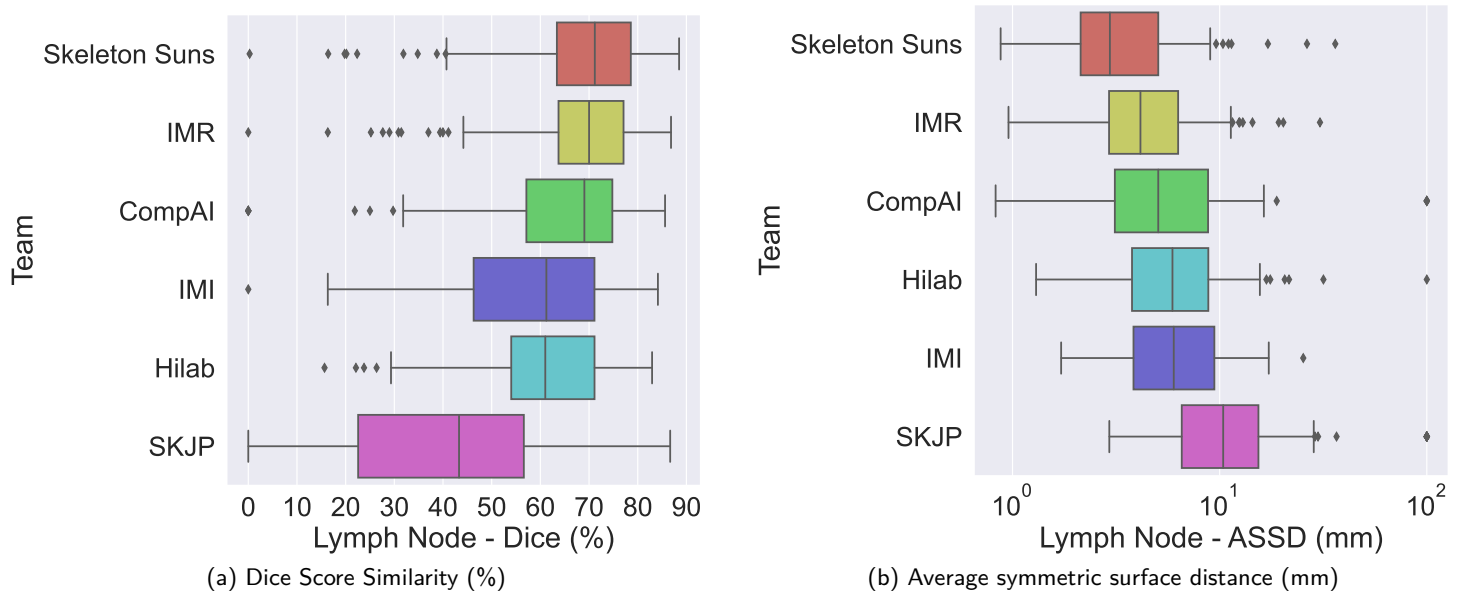


Figure 2: Box plot of the participants' performance for lymph node segmentation in terms of (a) DSC and (b) ASSD.

lymph nodes associated with Hodgkin lymphoma are typically bulky and organized in conglomerates, making their segmentation more challenging.

6.3 Remarks about the ranking stability

Several design factors can influence challenge rankings, such as the test set used for validation and the aggregation method applied to these metrics (Maier-Hein et al., 2018). In this section, we analyze and visualize the stability of rankings with respect to these factors.

To evaluate ranking stability in the context of sampling variability, we adopted the approach described by Wiesenfarth et al. (2021) and used for the crossMoDA challenge (Dorent et al., 2023). Following their guidelines, we performed bootstrapping with 1,000 samples to examine the uncertainty and stability of our proposed ranking scheme. The ranking strategy was applied repeatedly to each bootstrap sample. Kendall's τ was used to quantify

the agreement between the original challenge ranking and the rankings derived from the bootstrap samples, yielding values between -1 (indicating reverse ranking order) and 1 (indicating identical ranking order). The median [IQR] Kendall's τ was 1 [0.87 - 1], indicating excellent stability of the ranking scheme. Figure 5 presents a blob plot of the bootstrap rankings, confirming the excellent stability, with the winning team consistently ranked first across all bootstrap samples.

We also compared our ranking method with other common aggregation approaches. The main methods are:

- **Aggregate-then-rank:** Metric values are first aggregated (e.g., mean, median) across all test cases for each structure and metric, then ranks are computed for each team. Final ranking scores are derived from aggregating these ranks.
- **Rank-then-aggregate:** Ranks are computed for each test

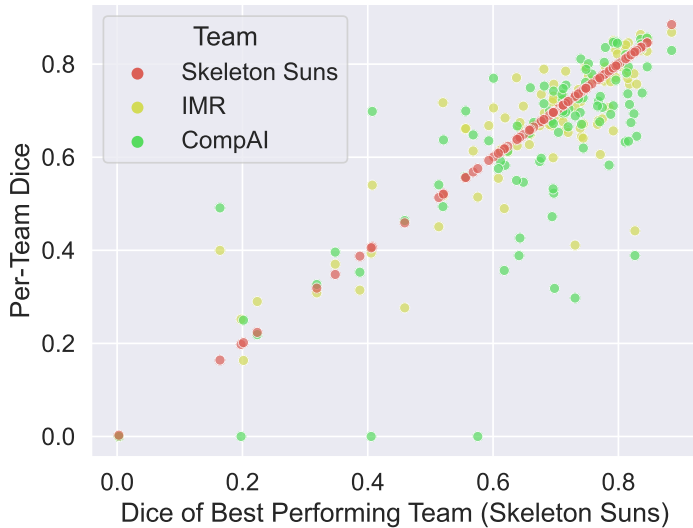


Figure 3: Relationship between the scores for each case between the first, second and third teams. Challenging cases are similar for each team.

case, metric, and structure, then aggregated (e.g., mean, median) to produce a final rank score for each algorithm.

Our method uses a rank-then-aggregate approach with the mean as the aggregation technique. We compared this to 1) a rank-then-aggregate approach using the median and 2) aggregate-then-rank approaches using either the mean or the median for metric aggregation. Line plots in Figure 6 illustrate the robustness of rankings across these methods. The ranks remained consistent across all ranking variations, with differences only occurring in cases of ties. Notably, Skeleton Suns was the top-performing team regardless of the ranking approach. This demonstrates that the challenge rankings are stable and can be interpreted with confidence.

7. Discussion and conclusion

In this study, we introduced the LNQ challenge in terms of experimental design, evaluation strategy, proposed methods, and final results. In this section, we discuss the main insights and limitations of the challenge.

7.1 Full supervision leads to higher performance

The top-ranked teams all utilized existing fully annotated datasets to train their frameworks. Notably, the top two teams first trained their models without the weakly annotated LNQ data, subsequently refining their models using partial annotations. TotalSegmentator was especially used to remove false positives. In contrast, methods that solely leveraged weak annotations underperformed compared to these fully supervised approaches. This underscores that full supervision still outperforms existing weakly supervised

approaches, even when using a smaller amount of training data.

Given that the proposed weakly supervised methods did not close the gap between full and weak supervision, we hope this challenge will continue to serve as a benchmark for developing new weakly supervised approaches.

Additionally, we believe this challenge will contribute to the improvement of existing models trained using full supervision, as the validation and testing datasets currently represent the largest publicly available fully annotated dataset (N=120 scans).

7.2 Limitations

This challenge was designed to benchmark new and existing weakly-supervised techniques for lymph node segmentation. In this section, we acknowledge some limitations.

First, there was a noticeable shift in the distribution of cancer types between the training, validation, and testing sets. Despite this gap, similar performance levels were observed across all cancer types, except for Hodgkin lymphoma. Importantly, the proportion of Hodgkin lymphoma cases was consistent across all three sets, suggesting that the lower performance for this type was due to the inherent difficulty in segmenting its bulky and conglomerated lymph nodes rather than the distribution gap.

The annotation process included selecting lymph nodes considered abnormal, with a short axis length larger than 1 cm. However, it is believed that some lymph nodes around this threshold may have been overlooked. Moreover, no post-processing was performed to remove those with a shorter axis. Some teams attempted to perform this post-processing, which led to degraded results on the test set.

In addition, the weakly annotated training dataset was created by completing volumetric segmentations of lymph nodes that had been selected for measurement during the TIMC clinical trial process, meaning that it may be biased towards the larger or more clinically significant lymph nodes compared to the nodes included in the fully annotated test and validation datasets.

Finally, a single-label map was used for all nodes. We acknowledge that it would have been beneficial to perform instance segmentation, where each node is individually segmented. This represents a potential area for future work.

7.3 Conclusion

The LNQ challenge was introduced to propose the first international benchmark for weakly-supervised image segmentation of lymph nodes in 3D CT scans, aimed at advancing the development of weak-supervised segmentation methods in the medical imaging community. By curating a new dataset and providing a standardized evalua-

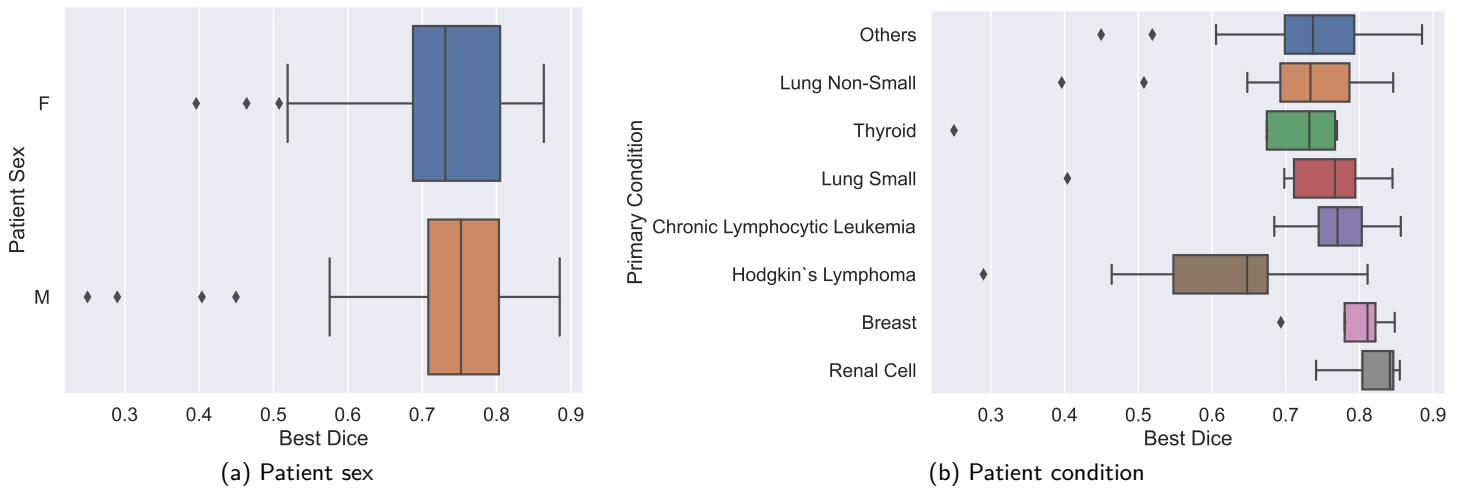


Figure 4: Box plot of the best performance per (a) patient sex and (b) primary condition.

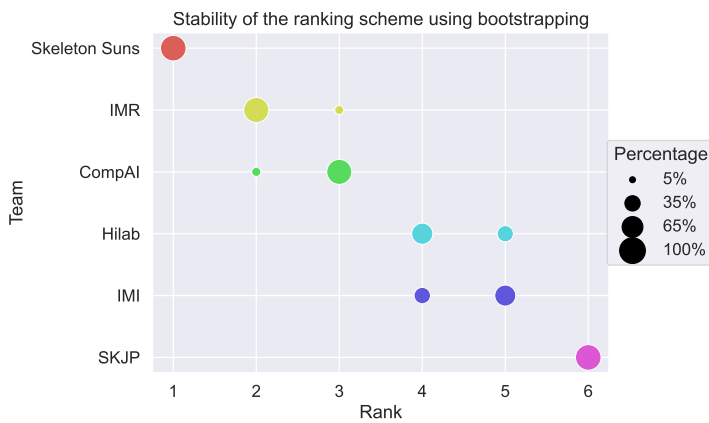


Figure 5: Stability of the proposed ranking scheme for 1000 bootstrap samples.

tion, we facilitated the comparison of different approaches and highlighted the current challenges and limitations of these weakly supervised techniques. Our findings indicate that fully-supervised methods, even when trained on smaller amounts of data, continue to outperform weakly-supervised approaches that leverage larger but partially annotated datasets. This highlights the ongoing need for high-quality, fully annotated data to achieve optimal segmentation performance. Nonetheless, the weakly-supervised methods showed promise, and we believe this challenge will encourage further innovation in this area. Overall, the LNQ challenge provides a valuable resource for the continued development and assessment of lymph node segmentation methods. The fully annotated validation and test sets, in particular, will serve as important assets for future research in lymph node quantification.

Acknowledgments

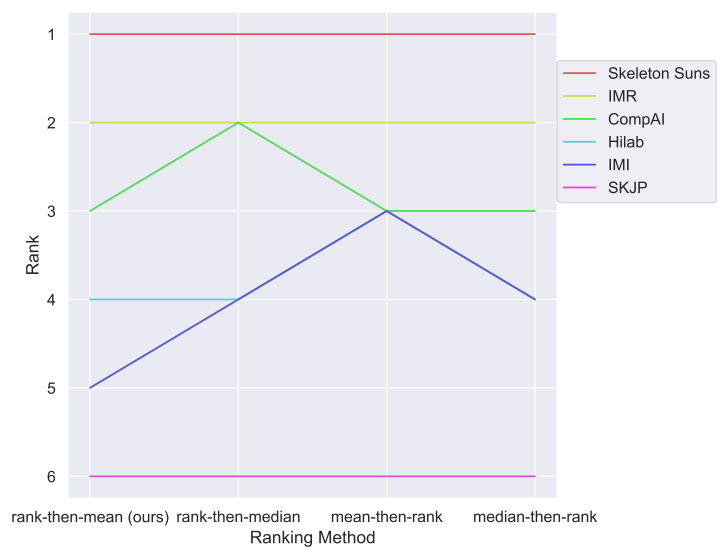


Figure 6: Line plots visualizing rankings robustness across different ranking methods for the brain task. Each algorithm is represented by one colored line. For each ranking method encoded on the x-axis, the height of the line represents the corresponding rank. The lowest rank of tied values is used for ties (equal scores for different teams)

This work was supported by the National Institutes of Health (R01EB032387, P41EB015902, P41EB028741, R01CA235589, 5P30CA006516, 1U24CA258511), National Cancer Data Ecosystem (Task Order No. 413 HHSN26110071 under Contract No. HHSN261201500003I) and EOSS5 and EOSS3 Diversity and Inclusion grants from Chan-Zuckerberg Initiative. R.D. received a Marie Skłodowska-Curie fellowship No 101154248 (project: SafeREG).

Ethical Standards

The work follows appropriate ethical standards in conducting research and writing the manuscript, following all applicable laws and regulations regarding treatment of animals or human subjects. Two IRB protocols were approved at Mass General Brigham covering de-identification of datasets assessed by the Tumor Imaging Metrics Core (TIMC) for the creation of a research data repository for downstream analysis (“Cancer Imaging Research Data Repository for the Lymph Node Quantification Project”, PI: Harris, Protocol Number: 2020P000211) and the annotations made during the clinical trials (Lymph Node Quantification (LNQ) Project”, PI: Kikinis, Protocol Number: 2020P003754).

Conflicts of Interest

G. Harris and E. Ziegler are co-founders of Yunu, Inc. G. Harris is also at the scientific advisory board of Fovia, Inc. The other 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.

Data availability

The data used in the challenge is currently available with registration on the LNQ grand-challenge website at <https://lnq2023.grand-challenge.org/> and will be released soon on TCIA as an unrestricted dataset.

References

- H. Aerts, Emmanuel Rios Velazquez, Ralph TH Leijenaar, Chintan Parmar, Patrick Grossmann, Sara Carvalho, Johan Bussink, René Monshouwer, Benjamin Haibe-Kains, Derek Rietveld, et al. Decoding tumour phenotype by noninvasive imaging using a quantitative radiomics approach. *Nature Communications*, 5(1):4006, 2014.
- Michela Antonelli, Annika Reinke, Spyridon Bakas, Keyvan Farahani, Annette Kopp-Schneider, Bennett A Landman, Geert Litjens, Bjoern Menze, Olaf Ronneberger, Ronald M Summers, et al. The medical segmentation decathlon. *Nature Communications*, 13(1):4128, 2022.
- Spyridon Bakas, Mauricio Reyes, Andras Jakab, Stefan Bauer, Markus Rempfler, Alessandro Crimi, Russell Takeshi Shinohara, Christoph Berger, Sung Min Ha, Martin Rozycki, et al. Identifying the best machine learning algorithms for brain tumor segmentation, progression assessment, and overall survival prediction in the BRATS challenge. *arXiv preprint arXiv:1811.02629*, 2018.
- Shaimaa Bakr, Olivier Gevaert, Sebastian Echegaray, Kelsey Ayers, Mu Zhou, Majid Shafiq, Hong Zheng, Jalen Anthony Benson, Weiruo Zhang, Ann NC Leung, et al. A radiogenomic dataset of non-small cell lung cancer. *Scientific Data*, 5(1):1–9, 2018.
- Adrian Barbu, Michael Suehling, Xun Xu, David Liu, S. Kevin Zhou, and Dorin Comaniciu. Automatic Detection and Segmentation of Lymph Nodes From CT Data. *IEEE Transactions on Medical Imaging*, 31(2):240–250, 2012.
- David Bouget, Arve Jørgensen, Gabriel Kiss, Haakon Olav Leira, and Thomas Langø. Semantic segmentation and detection of mediastinal lymph nodes and anatomical structures in CT data for lung cancer staging. *International Journal of Computer Assisted Radiology and Surgery*, 14:977–986, 2019.
- David Bouget, André Pedersen, Johanna Vanel, Haakon O Leira, and Thomas Langø. Mediastinal lymph nodes segmentation using 3D convolutional neural network ensembles and anatomical priors guiding. *Computer Methods in Biomechanics and Biomedical Engineering: Imaging & Visualization*, 11(1):44–58, 2023.
- Jinzheng Cai, Youbao Tang, Le Lu, Adam P Harrison, Ke Yan, Jing Xiao, Lin Yang, and Ronald M Summers. Accurate Weakly-Supervised Deep Lesion Segmentation Using Large-Scale Clinical Annotations: Slice-Propagated 3D Mask Generation from 2D RECIST. In *Medical Image Computing and Computer Assisted Intervention—MICCAI 2018: 21st International Conference, Granada, Spain, September 16–20, 2018, Proceedings, Part IV 11*, pages 396–404. Springer, 2018.
- Yigit B Can, Krishna Chaitanya, Basil Mustafa, Lisa M Koch, Ender Konukoglu, and Christian F Baumgartner. Learning to segment medical images with scribble-supervision alone. In *Deep Learning in Medical Image Analysis and Multimodal Learning for Clinical Decision Support: 4th International Workshop, DLMIA 2018, and 8th International Workshop, ML-CDS 2018, Held in Conjunction with MICCAI 2018, Granada, Spain, September 20, 2018, Proceedings 4*, pages 236–244. Springer, 2018.
- M. Jorge Cardoso, Wenqi Li, Richard Brown, Nic Ma, Eric Kerfoot, Yiheng Wang, Benjamin Murrey, Andriy Myronenko, Can Zhao, Dong Yang, Vishwesh Nath, Yufan He, Ziyue Xu, Ali Hatamizadeh, Andriy Myronenko, Wentao Zhu, Yun Liu, Mingxin Zheng, Yucheng Tang, Isaac Yang, Michael Zephyr, Behrooz Hashemian, Sachidanand Alle, Mohammad Zalbagi Darestani, Charlie Budd, Marc Modat, Tom Vercauteren, Guotai Wang,

- Yiwen Li, Yipeng Hu, Yunguan Fu, Benjamin Gorman, Hans Johnson, Brad Genereaux, Barbaros S. Erdal, Vikash Gupta, Andres Diaz-Pinto, Andre Dourson, Lena Maier-Hein, Paul F. Jaeger, Michael Baumgartner, Jayashree Kalpathy-Cramer, Mona Flores, Justin Kirby, Lee A. D. Cooper, Holger R. Roth, Daguang Xu, David Bericat, Ralf Floca, S. Kevin Zhou, Haris Shuaib, Keyvan Farahani, Klaus H. Maier-Hein, Stephen Aylward, Prerna Dogra, Sebastien Ourselin, and Andrew Feng. MONAI: An open-source framework for deep learning in healthcare, 2022.
- Reuben Dorent, Samuel Joutard, Jonathan Shapey, Sotirios Bisdas, Neil Kitchen, Robert Bradford, Shaheel Saeed, Marc Modat, Sébastien Ourselin, and Tom Vercauteren. Scribble-based domain adaptation via co-segmentation. In *Medical Image Computing and Computer Assisted Intervention—MICCAI 2020: 23rd International Conference, Lima, Peru, October 4–8, 2020, Proceedings, Part I 23*, pages 479–489. Springer, 2020.
- Reuben Dorent, Thomas Booth, Wenqi Li, Carole H Sudre, Sina Kafiabadi, Jorge Cardoso, Sebastien Ourselin, and Tom Vercauteren. Learning joint segmentation of tissues and brain lesions from task-specific hetero-modal domain-shifted datasets. *Medical Image Analysis*, 67: 101862, 2021a.
- Reuben Dorent, Samuel Joutard, Jonathan Shapey, Aaron Kujawa, Marc Modat, Sébastien Ourselin, and Tom Vercauteren. Inter extreme points geodesics for end-to-end weakly supervised image segmentation. In *Medical Image Computing and Computer Assisted Intervention—MICCAI 2021: 24th International Conference, Strasbourg, France, September 27–October 1, 2021, Proceedings, Part II 24*, pages 615–624. Springer, 2021b.
- Reuben Dorent, Aaron Kujawa, Marina Ivory, Spyridon Bakas, Nicola Rieke, Samuel Joutard, Ben Glocker, Jorge Cardoso, Marc Modat, Kayhan Batmanghelich, et al. CrossMoDA 2021 challenge: Benchmark of cross-modality domain adaptation techniques for vestibular schwannoma and cochlea segmentation. *Medical Image Analysis*, 83:102628, 2023.
- Elizabeth A Eisenhauer, Patrick Therasse, Jan Bogaerts, Lawrence H Schwartz, Danielle Sargent, Robert Ford, Janet Dancey, S Arbut, Steve Gwyther, Margaret Mooney, et al. New response evaluation criteria in solid tumours: revised RECIST guideline (version 1.1). *European Journal of Cancer*, 45(2):228–247, 2009.
- Sofija Engelson, Jan Ehrhardt, Timo Kepp, Joshua Niemeijer, and Heinz Handels. Lnq challenge 2023: Learning mediastinal lymph node segmentation with a probabilistic lymph node atlas. *Machine Learning for Biomedical Imaging*, 2:817–833, 2024a. ISSN 2766-905X. URL <https://melba-journal.org/2024:009>.
- Sofija Engelson, Jan Ehrhardt, Joshua Niemeijer, Stefanie Schierholz, Lennart Berkel, Malte Maria Elser, Yannic Sieren, and Heinz Handels. Comparison of Anatomical Priors for Learning-based Neural Network Guidance for Mediastinal Lymph Node Segmentation. In Weijie Chen and Susan M. Astley, editors, *Medical Imaging 2024: Computer-Aided Diagnosis*, volume 12927, page 1292719. International Society for Optics and Photonics, SPIE, 2024b.
- Andriy Fedorov, Reinhard Beichel, Jayashree Kalpathy-Cramer, Julien Finet, Jean-Christophe Fillion-Robin, Sonia Pujol, Christian Bauer, Dominique Jennings, Fiona Fennessy, Milan Sonka, et al. 3D Slicer as an image computing platform for the Quantitative Imaging Network. *Magnetic resonance imaging*, 30(9):1323–1341, 2012.
- Johannes Feulner, S Kevin Zhou, Matthias Hammon, Joachim Hornegger, and Dorin Comaniciu. Lymph node detection and segmentation in chest CT data using discriminative learning and a spatial prior. *Medical Image Analysis*, 17(2):254–270, 2013.
- Stefan M. Fischer, Johannes Kiechle, Daniel M. Lang, Jan C. Peeken, and Julia A. Schnabel. Mask the unknown: Assessing different strategies to handle weak annotations in the miccai2023 mediastinal lymph node quantification challenge. *Machine Learning for Biomedical Imaging*, 2:798–816, 2024. ISSN 2766-905X.
- Johannes Hofmanninger, Florian Prayer, Jeanny Pan, Sebastian Röhrich, Helmut Prosch, and Georg Langs. Automatic lung segmentation in routine imaging is primarily a data diversity problem, not a methodology problem. *European Radiology Experimental*, 4:1–13, 2020.
- Tagwa Idris, Bhanusupriya Somarouthu, Heather Jacene, Ann LaCasce, Erik Ziegler, Steve Pieper, Roya Khajavi, Reuben Dorent, Sonia Pujol, Ron Kikinis, and Gordon Harris. Mediastinal Lymph Node Quantification (LNQ): Segmentation of Heterogeneous CT Data. *The Cancer Imaging Archive*, 2024.
- Fabian Isensee and Klaus H Maier-Hein. An attempt at beating the 3D U-Net. *arXiv preprint arXiv:1908.02182*, 2019.
- Fabian Isensee, Paul F Jaeger, Simon AA Kohl, Jens Petersen, and Klaus H Maier-Hein. nnU-Net: a self-configuring method for deep learning-based biomedical

- image segmentation. *Nature methods*, 18(2):203–211, 2021.
- A. Emre Kavur, N. Sinem Gezer, Mustafa Barış, Sinem Aslan, Pierre-Henri Conze, Vladimir Groza, Duc Duy Pham, Soumick Chatterjee, Philipp Ernst, Savaş Özkan, Bora Baydar, Dmitry Lachinov, Shuo Han, Josef Pauli, Fabian Isensee, Matthias Perkonigg, Rachana Sathish, Ronnie Rajan, Debodoot Sheet, Gurbandurdy Dovletov, Oliver Speck, Andreas Nürnberger, Klaus H. Maier-Hein, Gözde Bozdağı Akar, Gözde Ünal, Oğuz Dicle, and M. Alper Selver. CHAOS Challenge - combined (CT-MR) healthy abdominal organ segmentation. *Medical Image Analysis*, 69:101950, 2021. ISSN 1361-8415.
- Hoel Kervadec, Jose Dolz, Shanshan Wang, Eric Granger, and Ismail Ben Ayed. Bounding boxes for weakly supervised segmentation: Global constraints get close to full supervision. In Tal Arbel, Ismail Ben Ayed, Marleen de Bruijne, Maxime Descoteaux, Herve Lombaert, and Christopher Pal, editors, *Proceedings of the Third Conference on Medical Imaging with Deep Learning*, volume 121 of *Proceedings of Machine Learning Research*, pages 365–381. PMLR, 06–08 Jul 2020.
- Roya Khajavi, Steve Pieper, Erik Ziegler, Tagwa Idris, Reuben Dorent, Bhanusupriya Somarouthu, Sonia Pujol, Ann LaCasce, Heather Jacene, Gordon Harris, and Ron Kikinis. Mediastinal Lymph Node Quantification (LNQ): Segmentation of Heterogeneous CT Data. *Zenodo*, April 2023. .
- Hyeonsoo Lee and Won-Ki Jeong. Scribble2label: Scribble-supervised cell segmentation via self-generating pseudo-labels with consistency. In *Medical Image Computing and Computer Assisted Intervention–MICCAI 2020: 23rd International Conference, Lima, Peru, October 4–8, 2020, Proceedings, Part I 23*, pages 14–23. Springer, 2020.
- Zhe Li and Yong Xia. Deep reinforcement learning for weakly-supervised lymph node segmentation in CT images. *IEEE Journal of Biomedical and Health Informatics*, 25(3):774–783, 2020.
- Lena Maier-Hein, Matthias Eisenmann, Annika Reinke, Sinan Onogur, Marko Stankovic, Patrick Scholz, Tal Arbel, Hrvoje Bogunovic, Andrew P. Bradley, Aaron Carass, Carolin Feldmann, Alejandro F. Frangi, Peter M. Full, Bram van Ginneken, Allan Hanbury, Katrin Honauer, Michal Kozubek, Bennett A. Landman, Keno März, Oskar Maier, Klaus Maier-Hein, Bjoern H. Menze, Henning Müller, Peter F. Neher, Wiro Niessen, Nasir Rajpoot, Gregory C. Sharp, Korsuk Sirinukunwattana, Stefanie Speidel, Christian Stock, Danail Stoyanov, Abdel Aziz Taha, Fons van der Sommen, Ching-Wei Wang, Marc-André Weber, Guoyan Zheng, Pierre Jannin, and Annette Kopp-Schneider. Why rankings of biomedical image analysis competitions should be interpreted with care. *Nature Communications*, 9(1):5217, Dec 2018. ISSN 2041-1723.
- Lena Maier-Hein, Annika Reinke, Michal Kozubek, Anne L. Martel, Tal Arbel, Matthias Eisenmann, Allan Hanbury, Pierre Jannin, Henning Müller, Sinan Onogur, Julio Saez-Rodriguez, Bram van Ginneken, Annette Kopp-Schneider, and Bennett A. Landman. BIAS: Transparent reporting of biomedical image analysis challenges. *Medical Image Analysis*, 66:101796, 2020. ISSN 1361-8415.
- Alireza Mehrtash, Erik Ziegler, Tagwa Idris, Bhanusupriya Somarouthu, Trinity Urban, Ann S LaCasce, Heather Jacene, Annick D Van Den Abbeele, Steve Pieper, Gordon Harris, et al. Evaluation of mediastinal lymph node segmentation of heterogeneous CT data with full and weak supervision. *Computerized Medical Imaging and Graphics*, 111:102312, 2024.
- Fausto Milletari, Nassir Navab, and Seyed-Ahmad Ahmadi. V-net: Fully convolutional neural networks for volumetric medical image segmentation. In *2016 fourth international conference on 3D vision (3DV)*, pages 565–571. IEEE, 2016.
- Cheng Ouyang, Chen Chen, Surui Li, Zeju Li, Chen Qin, Wenjia Bai, and Daniel Rueckert. Causality-Inspired Single-Source Domain Generalization for Medical Image Segmentation. *IEEE Transactions on Medical Imaging*, 42:1095–1106, 2021.
- Xi Ouyang, Zhong Xue, Yiqiang Zhan, Xiang Sean Zhou, Qingfeng Wang, Ying Zhou, Qian Wang, and Jie-Zhi Cheng. Weakly supervised segmentation framework with uncertainty: A study on pneumothorax segmentation in chest X-ray. In *Medical Image Computing and Computer Assisted Intervention–MICCAI 2019: 22nd International Conference, Shenzhen, China, October 13–17, 2019, Proceedings, Part VI 22*, pages 613–621. Springer, 2019.
- Steve Pieper, Michael Halle, and Ron Kikinis. 3D Slicer. In *2004 2nd IEEE international symposium on biomedical imaging: nano to macro (IEEE Cat No. 04EX821)*, pages 632–635. IEEE, 2004.
- Holger Roth, Le Lu, Ari Seff, Kevin M Cherry, Joanne Hoffman, Shijun Wang, Jiamin Liu, Evrim Turkbey, and Ronald M. Summers. A new 2.5 D representation for lymph node detection in CT (CT Lymph Nodes). *The Cancer Imaging Archive*, 2015.

- Roth, Holger and Zhang, Ling and Yang, Dong and Milletari, Fausto and Xu, Ziyue and Wang, Xiaosong and Xu, Daguang. Weakly supervised segmentation from extreme points. In *Large-Scale Annotation of Biomedical Data and Expert Label Synthesis and Hardware Aware Learning for Medical Imaging and Computer Assisted Intervention: International Workshops, LABELS 2019, HAL-MICCAI 2019, and CuRIOUS 2019, Held in Conjunction with MICCAI 2019, Shenzhen, China, October 13 and 17, 2019, Proceedings 4*, pages 42–50. Springer, 2019.
- Seyed Sadegh Mohseni Salehi, Deniz Erdogmus, and Ali Gholipour. Tversky loss function for image segmentation using 3D fully convolutional deep networks. In *International workshop on machine learning in medical imaging*, pages 379–387. Springer, 2017.
- Sarah Schuegger. MIC-DKFZ BodyPartRegression. *Zenodo*, August 2021. .
- Liza J. Stapleford, Joshua D. Lawson, Charles Perkins, Scott Edelman, Lawrence Davis, Mark W. McDonald, Anthony Waller, Eduard Schreiber, and Tim Fox. Evaluation of Automatic Atlas-Based Lymph Node Segmentation for Head-and-Neck Cancer. *International Journal of Radiation Oncology - Biology - Physics*, 77(3):959–966, 2010. ISSN 0360-3016.
- Mingxing Tan and Quoc Le. Efficientnet: Rethinking model scaling for convolutional neural networks. In *International Conference on Machine Learning*, pages 6105–6114. PMLR, 2019.
- Yongqiang Tan, Lin Lu, Apurva Bonde, Deling Wang, Jing Qi, Lawrence H. Schwartz, and Binsheng Zhao. Lymph node segmentation by dynamic programming and active contours. *Medical Physics*, 45(5):2054–2062, 2018.
- Litingyu Wang, Yijie Qu, Xiangde Luo, Wenjun Liao, Shichuan Zhang, and Guotai Wang. Weakly supervised lymph nodes segmentation based on partial instance annotations with pre-trained dual-branch network and pseudo label learning. *Machine Learning for Biomedical Imaging*, 2:1030–1047, 2024. ISSN 2766-905X. . URL <https://melba-journal.org/2024:017>.
- Yisen Wang, Xingjun Ma, Zaiyi Chen, Yuan Luo, Jinfeng Yi, and James Bailey. Symmetric cross entropy for robust learning with noisy labels. In *IEEE/CVF International Conference on Computer Vision*, pages 322–330, 2019.
- Jakob Wasserthal, Hanns-Christian Breit, Manfred T. Meyer, Maurice Pradella, Daniel Hinck, Alexander W. Sauter, Tobias Heye, Daniel T. Boll, Joshy Cyriac, Shan Yang, Michael Bach, and Martin Segeth. TotalSegmentator: Robust Segmentation of 104 Anatomic Structures in CT Images. *Radiology: Artificial Intelligence*, 5(5):e230024, 2023.
- L. Wee, HJL Aerts, P Kalendralis, and A Dekker. Data From NSCLC-Radiomics-Interobserver1 [Data set]. *The Cancer Imaging Archive*, 10, 2019.
- Manuel Wiesenfarth, Annika Reinke, Bennett A. Landman, Matthias Eisenmann, Laura Aguilera Saiz, M. Jorge Cardoso, Lena Maier-Hein, and Annette Kopp-Schneider. Methods and open-source toolkit for analyzing and visualizing challenge results. *Scientific Reports*, 11(1):2369, Jan 2021. ISSN 2045-2322.
- Hanxiao Zhang, Minghui Zhang, Yun Gu, and Guang-Zhong Yang. Deep anatomy learning for lung airway and artery-vein modeling with contrast-enhanced CT synthesis. *International Journal of Computer Assisted Radiology and Surgery*, 18(7):1287–1294, 2023.
- Ke Zhang and Xiahai Zhuang. Cyclemix: A holistic strategy for medical image segmentation from scribble supervision. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 11656–11665, 2022.
- Man Zhang, Yong Zhou, Jiaqi Zhao, Yiyun Man, Bing Liu, and Rui Yao. A survey of semi-and weakly supervised semantic segmentation of images. *Artificial Intelligence Review*, 53:4259–4288, 2020.
- Binsheng Zhao, Leonard P James, Chaya S Moskowitz, Pingzhen Guo, Michelle S Ginsberg, Robert A Lefkowitz, Yilin Qin, Gregory J Riely, Mark G Kris, and Lawrence H Schwartz. Evaluating variability in tumor measurements from same-day repeat CT scans of patients with non-small cell lung cancer. *Radiology*, 252(1):263–272, 2009.
- Zhedong Zheng and Yi Yang. Rectifying pseudo label learning via uncertainty estimation for domain adaptive semantic segmentation. *International Journal of Computer Vision*, 129(4):1106–1120, 2021.
- Zongwei Zhou, Vatsal Sodha, Jiaxuan Pang, Michael B Gotway, and Jianming Liang. Models genesis. *Medical Image Analysis*, 67:101840, 2021.