

Automatic Lymph Node Cluster Segmentation using Holistically-Nested Neural Networks and Structured Optimization in CT Images

Isabella Nogues¹, Le Lu¹, Xiaosong Wang¹, Holger Roth¹, Gedas Bertasius², Nathan Lay¹, Jianbo Shi², Yohannes Tsehay¹, and Ronald M. Summers¹

¹Imaging Biomarkers and Computer-Aided Diagnosis Laboratory, Radiology and Imaging Sciences, National Institutes of Health Clinical Center, Bethesda, MD 20892-1182, USA

²University of Pennsylvania, Philadelphia, PA 19104, USA

Abstract. Lymph node segmentation is an important yet challenging problem in medical image analysis. The presence of enlarged lymph nodes (LNs) signals the onset or progression of a malignant disease or infection. In the thoracoabdominal (TA) body region, neighboring enlarged LNs often spatially collapse into “swollen” lymph node clusters (LNCs) (up to 9 LNs in our dataset). Accurate segmentation of TA LNCs is complexified by the noticeably poor intensity and texture contrast among neighboring LNs and surrounding tissues, and has not been addressed in previous work. This paper presents a novel approach to TA LNC segmentation that combines holistically-nested neural networks (HNNs) and structured optimization (SO). Two HNNs, built upon recent fully convolutional neural networks (FCNs) and deeply supervised networks (DSNs), are trained to learn the LNC appearance (HNN-A) or contour (HNN-C) probabilistic output maps, respectively. HNN first produces the class label maps with the same resolution as the input image, like FCN. Afterwards, HNN predictions for LNC appearance and contour cues are formulated into the unary and pairwise terms of conditional random fields (CRFs), which are subsequently solved using one of three different SO methods: dense CRF, graph cuts, and boundary neural fields (BNF). BNF yields the highest quantitative results. Its mean Dice coefficient between segmented and ground truth LN volumes is $82.1\% \pm 9.6\%$, compared to $73.0\% \pm 17.6\%$ for HNN-A alone. The LNC relative volume (mm^3) difference is $13.7\% \pm 13.1\%$, a promising result for the development of LN imaging biomarkers based on volumetric measurements.

1 Introduction

Lymph node (LN) segmentation and volume measurement play a crucial role in important medical imaging based diagnosis tasks, such as quantitatively evaluating disease progression or the effectiveness of a given treatment or therapy. Enlarged LNs, defined by the widely observed RECIST criterion [14] to have a short axis diameter ≥ 10 mm on an axial computed tomography (CT) slice, signal the onset or progression of a malignant disease or an infection. Often performed manually, LN segmentation is highly complex, tedious and time consuming. Previous methods for automatic LN segmentation in CT images fall under several categories, including atlas registration and label




Fig. 1. CT images of thoracoabdominal lymph node clusters with annotated (red) boundaries.

fusion [17], 3D deformable surface shape model [6] and statistical 3D image feature learning [1, 7], respectively. This paper addresses and solves a novel problem: lymph node cluster (LNC) segmentation in the thoracoabdominal (TA) region. LN volumes are subsequently predicted from our segmentation results.

In CT, the TA region exhibits exceptionally poor intensity and texture contrast among neighboring LNs and between LNs and their surrounding tissues. Furthermore, LNs in this region often appear in clusters. Weak intensity contrast renders the boundaries of distinct agglomerated LNs ambiguous (Fig. 1). Existing fully-automated methods have been applied to the more contrast-distinctive axillary and pelvic regions [1], as well as the head-and-neck section [6, 17]. This paper presents a fully-automated method for TA LNC segmentation. More importantly, the segmentation task is formulated as a flexible, bottom-up image binary classification problem that can be effectively solved using deep convolutional neural networks (CNN) and graph-based structured optimization and inference. Our bottom-up approach can easily handle all variations in LNC size and spatial configuration. By contrast, the top-down model-fitting based approaches [1, 6, 7, 17] may struggle to localize and segment each individual LN. LN volume is a more robust metric than short-axis diameter length, which is susceptible to high inter-observer variability and human error. Furthermore, our proposed method is well-suited for measuring agglomerated LNs, whose ambiguous boundaries compromise the accuracy of diameter measurement.

This paper addresses a clinically relevant and challenging problem: automatic segmentation and volume measurement of LNCs in the TA region. A publicly available dataset¹ containing 171 3D CT volumes (with manually-annotated LN segmentation masks) from this region [15] is used. The method in this paper integrates HNN learning with structured optimization (SO). Furthermore, it yields remarkable quantitative results. The mean Dice similarity coefficient (DSC) between predicted and segmented LN volumes is $82.1\% \pm 9.6\%$ for HNN-A + boundary neural fields (BNF), and $73.0\% \pm 17.6\%$ for HNN-A alone. The relative volume measurement error is $13.7\% \pm 13.1\%$ for HNN-A + BNF and $32.16\% \pm 36.28\%$ for HNN-A.

2 Methods

Our segmentation framework comprises two stages: holistically-nested neural network (HNN) training/inference and structured optimization (SO). 1) Two HNNs, designed

¹ <https://wiki.cancerimagingarchive.net/display/Public/CT+Lymph+Nodes>.

in [18], are trained on pairs of raw CT images from the TA region and their corresponding binary LN appearance (segmentation) or contour (boundary) masks. We denote them as HNN-A and HNN-C, respectively. The HNN merges the CNN frameworks of a fully convolutional neural network (FCN) [11] and a deeply-supervised network (DSN) [10]. The FCN component is an end-to-end holistic image training-prediction architecture: the output probability label map has the same dimension as the input image. The DSN component performs multi-scale feature learning, in which deep layer supervision informs and refines classification results at multiple convolutional stages. HNN’s multi-level contextual architecture and auxiliary cost functions (which assign pixel-wise label penalties) allow for capturing implicit, informative deep features to enhance the segmentation accuracy.

2) However, HNN’s large receptive fields and pooling layers potentially lead to segmentation outputs that are not precisely localized along the LN boundaries (Fig. 2). Hence, we implement and evaluate explicit Conditional Random Field (CRF) based structured optimization schemes [2, 3, 9] to further improve segmentation results. Particularly, we seek to optimize a CRF energy function that includes a unary term encoding LN area information (HNN-A) and a pairwise term representing the object-specific LN boundary discontinuities (learned by HNN-C via deep supervision). This pairwise energy function differs from the conventional intensity contrast-sensitive terms described in [3] and [9].

We note that the integration of boundary and appearance neural networks for automatic segmentation has also been newly exploited in [13]. Indeed, both this paper and [13] have been partially inspired by visual cognitive science discoveries and the computer vision literature (namely [12, 18]). However, the proposed decision integration representations, principles and implementations in each work differ significantly: [13] refines HNN predictions with superpixel and random forest aggregation, as opposed to graph-based structured optimization.

2.1 Holistically-Nested Neural Networks

The holistically-nested neural network (HNN) [18] was first proposed as an image-to-image solution to the long-standing edge detection problem using deep CNN. In this study, we empirically find that the HNN architecture, in addition to predicting object boundaries with high accuracy, is highly effective and efficient in predicting the overall object segmentation mask, due to its per-pixel label cost formulation. Therefore, we train two separate holistically-nested neural networks to learn the probabilistic label maps of the LN specific binary appearance mask (HNN-A) and contour mask (HNN-C) from raw TA CT images. The HNN-A prediction map provides the approximate location and shape of the target LNC, while the HNN-C prediction map renders LN boundary cues. By learning boundaries alone, HNN-C generates boundary information with refined spatial and contextual detail, relative to HNN-A. HNN-A and HNN-C results are combined in the Structured Optimization phase (cf. Section: Structured Optimization) to obtain more accurate pixel-wise label predictions.

HNN Training We adopt the CNN architecture in [18], which derives from a VGGNet model pre-trained on ImageNet [16]. The HNN contains five convolutional stages, with




Fig. 2. Left: A cropped CT image with a lymph node located in the center is shown with its associated HNN-A (**Middle**) and HNN-C (before non-maximum suppression) (**Right**) maps.

strides 1, 2, 4, 8, and 16, respectively, and different receptive field sizes, all nested in the VGGNet as in [18]. The HNN includes one side layer per convolutional stage, which is associated with an auxiliary classifier. The side outputs, generated by each side layer, are increasingly refined, as they gradually approach the ground truth. Finally, all side outputs are fed into a "weighted-fusion" layer, which generates a global probability map merging the information from all side output scales. During the training phase, HNN seeks to minimize the per-pixel cost of each network stage, by applying stochastic gradient descent to the global objective function

$$(\mathbf{W}, \mathbf{w}, \mathbf{h})^* = \operatorname{argmin}(\mathcal{L}_{\text{side}}(\mathbf{W}, \mathbf{w}) + \mathcal{L}_{\text{fuse}}(\mathbf{W}, \mathbf{w}, \mathbf{h})), \quad (1)$$

where $\mathcal{L}_{\text{side}}$ is the loss function computed at each side-output layer (i.e., auxiliary cost functions), and $\mathcal{L}_{\text{fuse}}$ is the cross-entropy distance between the ground truth and fusion layer output edge maps. The loss function $\mathcal{L}_{\text{side}}$ is a linear combination of the image-level loss functions $\ell_{\text{side}}^{(m)}(\mathbf{W}, \mathbf{w}^{(m)})$. The parameters \mathbf{W} correspond to the set of standard network parameters, and \mathbf{w} to the weights in each side-output layer's classifier. Due to the per-pixel cost setup [11, 18], HNN does not require a large number of training images to converge, and thus can be applied to small datasets.

HNN Testing During the testing phase, the network generates edge map predictions for each layer. The final unified output is a weighted average of all prediction maps, denoted by:

$$\hat{Y}_{\text{HED}} = \text{Average}(\hat{Y}_{\text{fuse}}, \hat{Y}_{\text{side}}^{(1)}, \dots, \hat{Y}_{\text{side}}^{(5)}) \quad (2)$$

HNN is highly efficient, requiring a mere 0.4 second per image in the feed-forward testing. Further details on the HNN architecture and training protocol are provided in [18].

2.2 Structured Optimization

Although HNN can be considered a state-of-the-art, image-to-image, semantic pixel-wise labeling method, it tends to produce imprecise segmentations, like many deep

CNN models. Indeed, its large receptive fields and many pooling layers compromise clarity and spatial resolution in the deep layers. Therefore, we exploit explicit structured optimization techniques to refine HNN-A’s segmentation results. We select a conditional random field (CRF) optimization framework, as it is well-suited for integrating LN predictions with LN boundary cues. As in [2], the boundary cues serve to improve segmentation coherence and object localization. The global objective function is

$$E(\mathbf{x}) = \sum_i \psi(x_i) + \sum_{i < j} \psi(x_i, x_j). \quad (3)$$

For a given target CT image, the unary potential $\psi(x_i)$ is computed independently for each pixel in the corresponding HNN-A prediction, while the pairwise potential $\psi(x_i, x_j)$ is a function of the corresponding HNN-C prediction. Three structured optimization representations and methods are described and evaluated: dense CRF, graph cuts, and boundary neural fields.

Under all three techniques, segmentation is cast as a binary classification problem. A graph representation for the original CT image is provided, in which vertices correspond to image pixels, and edges to inter-pixel connections. A boundary strength-based affinity function, defined in [2], for distinct pixels i and j is given by

$$w_{ij} = \exp\left(\frac{-M_{ij}}{\sigma}\right), \quad (4)$$

where M_{ij} is the magnitude of the strongest LN boundary intersecting $\{i, j\}$, and σ is a smoothing hyper-parameter. The boundary strength map is obtained by performing non-maximum suppression on an HNN-C prediction [18]. This boundary-based affinity function better refines the segmentation than would a standard image intensity gradient-based function, as intensity information is highly ambiguous in TA CT images. We set the degree of pixel i to $d_i = \sum_{i \neq j}^N w_{ij}$, where N is the total number of pixels.

Dense Conditional Random Field A dense conditional random field (dCRF) representation is applied on top of the HNN-A predictions, following the framework in [5, 9]. The CT intensity contrast-sensitive pixel affinities for all possible pairs of image pixels are adopted, as described in [5]. Finally, the dCRF solver designed in [9] is utilized, as a variation of distributed message passing. Further experimental information is provided in [9].

Graph Cuts A graph cuts (GC) algorithm, namely the minimum-cut/maximum-flow algorithm described in [4], is applied to the segmentation problem, in order to compute the optimal bi-partition of the pixels in the image-defined graph. We optimize an energy function whose unary term is the negative log-likelihood of the HNN-A LN segmentation probability value per pixel and pairwise term is defined using the boundary-based affinity function presented above within a 20×20 neighborhood for each pixel location.

Boundary Neural Fields The LN mask (HNN-A) and boundary (HNN-C) predictions are integrated into a matrix model. We optimize the global energy function:

$$\mathbf{X}^* = \underset{\mathbf{X}}{\operatorname{argmin}} \frac{\mu}{2} \mathbf{D}(\mathbf{X} - \mathbf{D}^{-1}\mathbf{f})^T (\mathbf{X} - \mathbf{D}^{-1}\mathbf{f}) + \frac{1}{2} \mathbf{X}^T (\mathbf{D} - \mathbf{W}) \mathbf{X}, \quad (5)$$

where \mathbf{X}^* is a $N \times 1$ vector representing an optimal continuous *label assignment* for a vectorized input image (with N pixels), \mathbf{D} is the $N \times N$ diagonal degree matrix, \mathbf{W} is the $N \times N$ pairwise affinity matrix, and \mathbf{f} is a $N \times 1$ vector containing the HNN-A prediction values. Each diagonal entry $d_{i,i}$ is set to d_i , defined above. \mathbf{W} is a sparse weight matrix: the entries w_{ij} are computed only for pixel pairs i, j belonging to the same 20×20 pixel neighborhood. (via Eq. 4 [2]).

The unary energy attempts to find a segmentation assignment \mathbf{X} that deviates little from the HNN-A output. The assignment \mathbf{X} is weighted by \mathbf{D} , in order to assign larger unary costs to pixels with many similar neighbors. By contrast, the pairwise energy minimizes the cost assigned to such pixels by weighting the squared distances between segmentation assignments of similar pixel pairs ij by their affinity w_{ij} . To balance the unary and pairwise contributions, the unary term is weighted by the hyperparameter μ . As in [2], μ is set to 0.025 in all experiments. The optimal segmentation is given by the closed-form solution:

$$\mathbf{X}^* = (\mathbf{D} - \alpha \mathbf{W})^{-1} \beta \mathbf{f}, \quad (6)$$

where $\alpha = \frac{1}{1+\mu}$ and $\beta = \frac{\mu}{1+\mu}$.

3 Results and Discussion

3.1 Dataset Creation

The dataset used in this paper derives from 84 abdominal and 87 mediastinal 3D CT scans ($512 \times 512 \times 512$ voxels) (publicly available from [15]). We spatially group the ground truth binary LN masks in 3D to form clusters with a linking distance constraint. All LN clusters, with 32 pixel padding in each direction, are subsequently cropped, resulting in 1~7 such subvolume regions (with dimensions ranging from $77 \times 76 \times 79$ to $212 \times 235 \times 236$ voxels) per CT volume. All CT axial slices have been extracted from the portal venous phase with slice thickness 1 – 1.25 mm, and have been manually segmented by an expert radiologist. This yields a total of 39,361 images (16,268 images containing LN pixels) in 411 LN clusters (containing 395 abdominal and 295 mediastinal LNs). By extracting all LN boundaries from the ground truth segmentation masks, we obtain the corresponding LN contour maps to train HNN-C. Examples of LN CT image ground truth boundaries are shown in Fig. 1 and in the first row of Fig. 3.

3.2 Quantitative Analysis

Segmentation accuracies of the HNN-A alone and of the HNN-A followed by each of the three optimization methods (where HNN-C information is integrated) are evaluated. The volume-wise means and standard deviations (std.) are computed for three evaluations metrics: Dice similarity coefficient (DSC), Intersection over Union (IoU), and Relative Volume Difference (RVD) (cm^3) between predicted and ground truth LNC 3D masks. The RVD indicates whether volume measurement is accurate enough to be used as a new imaging biomarker, in addition to the diameter-based RECIST criterion [14].

Our experiments are conducted under 4-fold cross-validation, with the dataset split at the patient level to create the training and validation folds. To generate binary segmentation results for HNN-A prediction maps alone, it is necessary to apply a simple




Fig. 3. Examples of LN CT image segmentation. **Top:** original CT images with ground truth (red) and BNF segmented (green) boundaries. **Center:** HNN-A LN probability maps. **Columns 1-3** (from left) depict successful segmentation results. **Column 4** presents an unsuccessful case. **Bottom:** Additional segmentation results.

thresholding. The optimal thresholding value that maximizes the mean DSC between the HNN-A predictions and ground truth LN masks is calibrated using the training folds. The HNN is run on the Caffe platform from [8], using a Nvidia Tesla K40 GPU. The system requires 5 hrs 40 min for training (30K iterations), and 7 min 43 sec for testing (computation of all contour and area maps). BNF is run on MATLAB 2013a (~ 9 hours). dCRF and GC are run using C++ (~ 3 hrs; ~ 12 hrs). The hyperparameters of dCRF and GC are optimized using a randomized search.

BNF yields the best quantitative segmentation results. Its mean and std per-volume DSC is $82.1 \pm 9.6\%$, above HNN-A's $73.0 \pm 17.6\%$, dCRF's $69.0 \pm 22.0\%$, and GC's $67.3 \pm 16.8\%$ (cf. Table 4). Additionally, it decreases HNN-A's mean RVD from 32.2% to 13.7%. Meanwhile, dCRF marginally decreases the RVD to 29.6%, and GC increases it to 86.5%. Fig. 3 illustrates three examples of LNC segmentation using BNF. The plots in Fig. 4 compare segmentation and ground truth LNC volume values (in cm^3).

Due to the HNN's deeply nested architecture and auxiliary loss functions on multi-scale side-output layers, the quality of HNN-A segmentation, after simple thresholding, is already high. dCRF's decline in performance relative to HNN-A may be partly attributed to its CT intensity contrast-sensitive pairwise CRF term, defined in [9]. While effective in natural images, this color-based function may compromise segmentation accuracy in TA LNC CT images, which exhibit low intensity contrast and significant clutter. The expression for the appearance kernel in the pairwise term of [9] implies that nearby pixels of similar intensity are likely to belong to the same class. Though highly relevant in natural images, this idea cannot be extended to TA CT images, in which LN and background pixels may have similar intensities. GC, however, uses a pairwise term defined by the HNN-C boundary cues. Its lower performance, then, may be attributed

to its usage of an L_1 norm for the CRF energy minimization. Unlike the L_2 norm, used by dCRF and BNF, the L_1 norm may yield unstable or multiple solutions, thus compromising the final segmentation accuracy. Like GC, BNF omits all intensity information from the energy function. However, it utilizes the boundary cues from HNN-C in both the unary and pairwise terms. Hence, one may infer that an emphasis on LN boundary information is crucial to overcome the complexity of TA LNC CT image segmentation. Additionally, the L_2 norm may increase the accuracy of the final segmentation result.

Furthermore, the GC and dCRF models include variable hyperparameters (defined in [3] and [9], respectively), which require extensive fine-tuning. In our experiments, fine-tuning has been performed via a randomized search using the training data. BNF, on the other hand, only uses a single parameter μ to balance the LN and boundary information, and can be solved with a closed-form solution as shown in Eq. 6.

Method	Evaluation Metric		
	Mean DSC (%)	Mean IoU (%)	Mean RVD (%)
HNN-A	73.0 ± 17.6	60.1 ± 18.8	32.2 ± 46.3
BNF	82.1 ± 9.6	70.6 ± 11.9	13.7 ± 13.1
dCRF	69.0 ± 22.0	56.2 ± 21.6	29.6 ± 45.4
GC	67.3 ± 16.8	53.0 ± 17.9	86.5 ± 107.6

Table 1. Evaluation of segmentation accuracy : HNN-A, BNF, dCRF, and GC

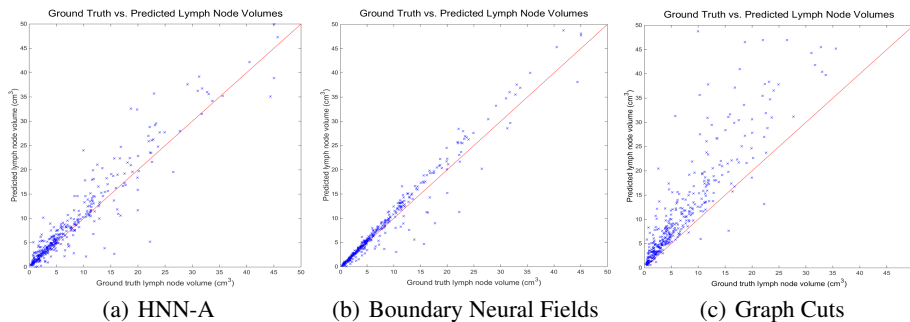


Fig. 4. Comparison between ground truth and predicted LN volumes.

Comparison to previous work: The proposed bottom-up LN segmentation scheme equals or surpasses previous top-down methods in performance, though applied to a more challenging body region. Indeed, [17] obtains a mean DSC of 73% on five patient head-and-neck CT scans and [6] shows the relative volumetric segmentation error ratios, ranging from 38.88% to 51.75% for five adult neck CT cases. For discriminative learning approaches, [7] only reports LN detection results on 54 chest CT datasets and no information on the LN segmentation accuracy. As one of the most comprehensive studies, [1] achieves a mean DSC of $80\% \pm 12.6\%$ for 308 Axillary LNs and $76\% \pm 12.7\%$ for 455 Pelvic+Abdominal LNs, from a dataset of 101 CT cases. Hence, one may infer that a bottom-up approach is most suitable for LN segmentation.

4 Conclusion

To solve a challenging problem with high clinical relevance – automatic segmentation and volume measurement of TA LNCs in CT images, our method integrates HNN learning in both LN appearance and contour channels and exploits different structured optimization methods. BNF (combining HNN-A and HNN-C via a sparse matrix representation) is the most accurate segmentation scheme. BNF’s mean RVD of $13.7 \pm 13.1\%$ is promising for the development of LN imaging biomarkers based on volumetric measurements, which may lay the groundwork for improved RECIST LN measurements.

Acknowledgments: This work was supported by the Intramural Research Program at the NIH Clinical Center.

References

1. A. Barbu, M. Suehling, X. Xu, D. Liu, S. K. Zhou, and D. Comaniciu. Automatic detection and segmentation of lymph nodes from ct data. *IEEE Trans. on Medical Imaging*, 2012.
2. G. Bertasius, J. Shi, and L. Torresani. Semantic segmentation with boundary neural fields. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016.
3. Y. Boykov and G. Funka-Lea. Graph cuts and efficient nd image segmentation. *IJCV*, 2006.
4. Y. Boykov and V. Kolmogorov. An experimental comparison of min-cut/max-flow algorithms for energy minimization in vision. In *IEEE Trans. on Pat. Ana. Mach. Intel.*, 2004.
5. L.-C. Chen, G. Papandreou, I. Kokkinos, K. Murphy, and A. L. Yuille. Semantic image segmentation with deep convolutional nets and fully connected crfs. In *ICLR*, 2015.
6. J. Dornheim, H. Seim, B. Preim, I. Hertel, and G. Strauss. Segmentation of neck lymph nodes in ct datasets with stable 3d mass-spring models. In *MICCAI*, pages 904–911, 2006.
7. J. Feulner, S. Zhou, M. Hammond, J. Hornegger, and D. Comaniciu. Lymph node detection and segmentation in chest ct data using discriminative learning and a spatial prior. In *Medical Image Analysis*, pages 254–270, 2013.
8. Y. Jia. Caffe: An open source convolutional architecture for fast feature embedding (2013). Available: <http://goo.gl/Fo9Y08>.
9. P. Krähenbühl and V. Koltun. Efficient inference in fully connected crfs with gaussian edge potentials. In *NIPS*, 2012.
10. C.-Y. Lee, S. Xie, P. Gallagher, Z. Zhang, and Z. Tu. Deeply-supervised nets. In *AISTATS*, 2015.
11. J. Long, E. Shelhamer, and T. Darrell. Fully convolutional networks for semantic segmentation. In *IEEE CVPR*, pages 3431–3440, 2015.
12. D. Martin, C. Fowlkes, and J. Malik. Learning to detect natural image boundaries using local brightness, color, and texture cues. In *IEEE Trans. on Pat. Ana. Mach. Intel.*, 2004.
13. H. Roth, L. Lu, A. Farag, A. Sohn, and R. Summers. Spatial aggregation of holistically-nested networks for automated pancreas segmentation. In *MICCAI (in print)*, 2016.
14. L. Schwartz, J. Bogaerts, R. Ford, L. Shankar, P. Therasse, S. Gwyther, and E. Eisenhauer. Evaluation of lymph nodes with recist 1.1. *Euro. J. of Cancer*, 45(2):261 – 267, 2009.
15. A. Seff, L. Lu, A. Barbu, H. Roth, H. Shin, and R. Summers. Leveraging mid-level semantic boundary cues for automated lymph node detection. In *MICCAI*, pages 2:53–61. 2015.
16. K. Simonyan and A. Zisserman. Very deep convolutional networks for large-scale image recognition. In *ICLR*, 2014.
17. L. Stapleford, J. Lawson, and et al. Evaluation of automatic atlas-based lymph node segmentation for head-and-neck cancer. In *Int. J. of Rad. Onc. Bio. Phys.*, 2010.
18. S. Xie and Z. Tu. Holistically-nested edge detection. In *IEEE ICCV*, pages 1395–1403, 2015.