PROCEEDINGS OF SPIE

SPIEDigitalLibrary.org/conference-proceedings-of-spie

Accurately identifying vertebral levels in large datasets

Elton, Daniel, Sandfort, Veit, Pickhardt, Perry, Summers, Ronald

Daniel C. Elton, Veit Sandfort, Perry J. Pickhardt M.D., Ronald M. Summers M.D., "Accurately identifying vertebral levels in large datasets," Proc. SPIE 11314, Medical Imaging 2020: Computer-Aided Diagnosis, 1131400 (16 March 2020); doi: 10.1117/12.2551247



Event: SPIE Medical Imaging, 2020, Houston, Texas, United States

Accurately identifying vertebral levels in large datasets

Daniel C. Elton^a, Veit Sandfort^a, Perry J. Pickhardt^b, and Ronald M. Summers^a

^aImaging Biomarkers and Computer-Aided Diagnosis Laboratory, Radiology and Imaging Sciences, National Institutes of Health Clinical Center, Bethesda, MD 20892-1182, USA ^bSchool of Medicine and Public Health, University of Wisconsin, Madison, USA

ABSTRACT

The vertebral levels of the spine provide a useful coordinate system when making measurements of plaque, muscle, fat, and bone mineral density. Correctly classifying vertebral levels with high accuracy is challenging due to the similar appearance of each vertebra, the curvature of the spine, and the possibility of anomalies such as fractured vertebrae, implants, lumbarization of the sacrum, and sacralization of L5. The goal of this work is to develop a system that can accurately and robustly identify the L1 level in large heterogeneous datasets. The first approach we study is using a 3D U-Net to segment the L1 vertebra directly using the entire scan volume to provide context. We also tested models for two class segmentation of L1 and T12 and a three class segmentation of L1, T12 and the rib attached to T12. By increasing the number of training examples to 249 scans using pseudo-segmentations from an in-house segmentation tool we were able to achieve 98% accuracy with respect to identifying the L1 vertebra, with an average error of 4.5 mm in the craniocaudal level. We next developed an algorithm which performs iterative instance segmentation and classification of the entire spine with a 3D U-Net. We found the instance based approach was able to yield better segmentations of nearly the entire spine, but had lower classification accuracy for L1.

Keywords: machine learning, deep learning, segmentation, instance segmentation, spine, vertebrae, U-Net

1. INTRODUCTION

The application of deep learning systems to radiology holds much promise for making medicine more quantitative.^{1,2} Deep learning based systems can be used opportunistically to gather valuable information from CT scans which is useful for screening for certain diseases and identifying risk factors. Examples of automated opportunistic measurements that have been demonstrated recently are bone mineral density,^{3,4} visceral fat,⁵ muscle,⁶ organ volume,⁷ plaque burden,⁸ and body composition analysis.⁹ If a patient has multiple CT scans performed, changes over time can be tracked. The resulting data is not only useful for patient care but also for scientific studies looking at health outcomes in large cohorts. However, in order for such systems to be useful consistent measurement is important. Being a rigid framework, the vertebrae of the spine provide a useful coordinate system for making measurements. Conventions have been established in this regard - for instance it is common to measure muscle content at L3,⁶ and bone mineral density at L1.⁴ Aortic plaque burden may be measured between the L1 and L4 levels.⁸

The thoracic vertebrae (T1-T12) and lumbar vertebrae (L1-L5) all have a similar appearance so classification based solely on a cropped bounding box around a given vertebra is difficult. Radiologists identify vertebrae based on their position relative to transitional vertebrae such as L1/T12 and L5, and we hypothesized that a deep learning system can learn to do something similar. However, classification and segmentation of the vertebrae is made challenging by the high variability of the spine. One source of variability is varying degrees of curvature due scoliosis/kyphosis or due to the patient not lying perfectly straight in the scanner. Another source of variability is the presence of lumbosacral transitional vertebrae, which occur in about 10% of patients.¹⁰ About 6% of patients have lumbarization of S1, and about 4% of patients have sacralization of L5.¹⁰ For this reason, L5 is not a reliable transitional vertebra to use as a reference landmark. Abnormalities also exist at the T12 and L1 vertebrae, such

Medical Imaging 2020: Computer-Aided Diagnosis, edited by Horst K. Hahn, Maciej A. Mazurowski, Proc. of SPIE Vol. 11314, 1131400 · © 2020 SPIE CCC code: 1605-7422/20/\$21 · doi: 10.1117/12.2551247

Further author information: Send correspondence to Daniel Elton (daniel.elton@nih.gov) or Ronald Summers rms@nih.gov

architecture details		N_{train}	N_{test}	acc.	citation(s)	
3D U-Net	Trained on spinal CT (33%) and low dose chest	60	30	94~%	Lessmann et al. 2019^{13}	
	CT (66%). Patch based iterative instance seg-					
	mentation.					
Two stage	Trained on spinal CT. Two stage network consist-	242	60	92~%	McCouat & Glocker 2019^{15}	
CNN.	ing of a 3D U-Net for detection of the spine and					
	a 2D U-Net for vertebral identification.					
3D CNN	Trained on spinal CT. Patch based 3D CNN used	242	60	92%	Liao et al., 2018^{16}	
+ FCNN	for initial labeling. Labels are refined using a					
+ biRNN	FCNN coupled to a bidirectional RNN.					
3D U-Net	Trained on various CT scans. Network performs	1112	112	84%	Yang et al., 2017^{17}	
	centroid placement and classification/labeling					
	only. Uses "deep supervision", message passing,					
	and sparsity regularization.					
FCNN +	Trained on spinal CT. A FCNN is used to regress	12	3	*	Janssens et al. 2016^{18}	
3D U-Net	a bound box around the lumbar region, and a					
	multiclass 3D U-Net is used to segment the 5 lum-					
	bar vertebrae.					

Table 1. Select prior works on vertebrae segmentation and identification using deep learning which report the highest accuracy. *Janssens et al. do not perform classification, but their segmentations of the 5 lumbar vertebrae had average Dice coefficients of 0.95 so we can surmise the correct vertebrae could be labeled nearly 100% of the time, at least within their small test set of 3 scans.

as the presence of 13th rib bearing vertebral bodies (incidence about $1\%^{11}$) or rudimentary/non-existent 12th ribs (incidence about $8\%^{12}$). In this work we choose to focus on using L1 as our reference vertebrae.

Several machine learning based approaches to segmenting and partitioning the spine have already been published. A full review is outside the scope of this paper, but we summarize the best-performing systems we found in table 1. We noticed that some prior studies (especially pre-2016) used small test sets (ie. 2-10 cases) which have a low probability of having abnormalities. One of the state of the art methods is a system developed by Lessmann, et al. which uses iterative instance segmentation with a 3D U-Net.^{13,14} Overall they achieved a 94% accuracy (correct identification rate) for classifying lumbar vertebrae. In this paper we compare direct segmentation of L1 with an iterative instance based approach which is similar to the approach developed by Lessmann et al.

2. METHODS

2.1 Dataset preparation

To construct a training set of L1 segmentations we initially hand segmented 52 scans using 3D Slicer. 37 of the scans were taken from the "CT Lymph Nodes" dataset^{19,20} of the Cancer Imaging Archive²¹ and 15 come from a set of scans taken for CT colonoscopy at National Naval Medical Center.²² To generate a larger training data set we used a previously developed automated bone mineral density (BMD) code,²³ to generate rough segmentations of the entire spine ("pseudomasks"). To partition the spine, the tool detects the spinal cord, performs a curved reformulation along the spinal cord centerline, and partitions the spine based on the bone density along the centerline. For segmentation the watershed algorithm is used, combined with directed graph search to reduce oversegmentation. The L1 level was identified manually in each case, and the segmentation at the L1 level was extracted from the pseudo-segmentations generated by the code. Manual levels were found for 225 cases from the aforementioned Naval dataset and 32 cases from the Cancer Imaging Archive, to generate L1 segmentations for a total of 257 cases.

To train the 3D U-Net for iterative instance segmentation the cervical, thoracic, and lumbar vertebra were labeled with integers between 1 and 24 and the labels were aligned so that the vertebra with label 20 matched the level that was determined manually to be L1. To generate the subvolumes, we first resampled all images to 1mm x 1mm and clip the images to Hounsfeld units between -100 and 2000. Training and inference are performed with patches of size 128x128x128, which is large enough to include an entire vertebra as well as at least one vertebra above and below. The training subvolumes are generated by centering a 128x128x128 box on each vertebra in the ground truth data, with minor random jitter of ± 2 cm in each direction. Data augmentation similar to the prior method was used (random XY flipping, random B-spline deformation, and random rotations). The center (target) vertebra is given the label '1' and all vertebra above and below are all given the labels '2' and '3' for use as the instance memory during training. The network can be trained with only the above vertebra in the instance memory ("top-down" mode), only the ones below ("bottom-up" mode), or by randomly selecting one or the other (a "bidrectional" network). 15-20% of the generated subvolumes were empty, which is essential in order to train the algorithm to ignore things that would otherwise be mistaken as vertebra, such as contrast or hip bone.

We first trained the instance-based network on the generated pseudomasks of the entire spine for the same 257 cases discussed above. We also trained with 2 scans from the 2019 MICCAI "VerSe" challenge, which consist of a wide variety of volumes with hand segmentations of all the vertebra present in the scan.²⁴ Since we are only interested in the lumbar and thoracic vertebrae we disregarded any cervical vertebra in the VerSe data. We also used 10 scans with manual segmentations of the lumbar vertebrae made available by Prof. Vrtovec at the Laboratory of Imaging Technologies at the University of Ljubljana.^{25, 26}

To test the accuracy of both approaches for predicting the L1 level we used 40 cases from a separate set of scans taken for CT colonography, where the L1 level was identified manually.

instance memory instance m

2.2 Identifying the L1 level

Figure 1. The 3D U-Net architectures used for identifying L1 (left) and for instance segmentation (right).

We tested segmenting the L1 vertebrae using a 3D U-Net^{27, 28} architecture (shown in fig. 1). The 3D U-Net is adapted from an architecture developed for organ segmentation²⁹ and is implemented in Pytorch. Each layer uses group normalization³⁰ (group size = 16), dropout (dropout rate = 0.3), and a leaky ReLU activation function. Images are re-sampled to 128x128x192 and clipped to between 100-1000 HU before normalization. For augmentation we used independent random flips along the x, y, and z directions, elastic deformations using the B-spline deformation function in SITK and random rotations between -20 and 20 degrees around a random axis. Training was performed on a 4 Titan X GPUs with a batch size of 4 and run for $\approx 100,000$ iterations.

2.3 Iterative instance segmentation

We hypothesized that performing a full spine segmentation would result in more robust identification of L1. The approach we took is a simplified and modified version of the iterative instance algorithm developed by Lessmann et al.^{13,14} To simplify their approach we do not include a third output branch to detect whether a vertebra is fully or partially visible, and we do not use a loss function which weights the surface of vertebra more heavily. We used our own custom 3D U-Net architecture which is shown in fig. 1.

The loss function we used is a weighted sum of the number of false positive voxels (FP) and number of false negative voxels (FN) plus a term for level regression error:

$$\mathcal{L} = \lambda_n \mathrm{FP}_{\mathrm{frac}} + \mathrm{FN}_{\mathrm{frac}} + \beta \left| p_L - t_L \right| \tag{1}$$

We did not use the Dice loss (1 - DICE = (FP + FN)/(2TP + FP + FN)) since it provides no training gradient when the box is empty (ie. when TP = 0). The factor λ_n in equation 1 puts a lower weight on false positives during the beginning of training. Since the number of background (empty) voxels is much greater than the number of foreground voxels in the training dataset, at the start of training false positives are a much larger problem than false negatives. If λ_n is too large then during training the model ends up just outputting an empty segmentation mask to minimize the loss function. We adjust λ_n as a function of the iteration number n starting with $\lambda_n = 0.05$ at n = 0 and increase it sigmoidally to $\lambda_n \approx 1$ when $n = n_{\text{max}}$. The factor β is a hand-tuned factor which balances the importance of two terms (we settled on $\beta = 1000$). We generated 30% of our training subvolumes to be empty, and set their "level" set to 0. We also tried removing the term $\beta |p_L - t_L|$ from the loss function (eqn.1) for empty boxes and (independently) adding an additional Dice loss term, but found that neither of these modifications led to an improvement. Training was performed on a single Titan Z GPU with a batch size of 6 and run for $\approx 50,000$ iterations.

During inference, the network can operate in either "top-down" or "bottom-up" mode - ie. starting from the top of the image and working down or vice versa. We found top-down to be easier since we can center the network on the spine, which is always present at the top of the image in our dataset. The network starts with a box centered on the average of the bone voxels in the first two slices, which is typically close to the spine in top-down mode. If less than 1 cm³ of vertebral bone is detected in this first subvolume, the algorithm then starts from the top corner of the images and begins scanning across the image with a stride of 64 the x and y directions. When more than a predefined threshold of bone is detected, the subvolume is moved so that it is centered on the detected bone. This process of centering is iterated until the center of the segmented bone changes by less than 2 mm between iterations or until a maximum of 5 iterations is reached. In order to make the algorithm robust, we had to force the algorithm only to move down when operating in top-down mode and only to move up when operating in bottom-up mode. At the end of the centering procedure the segmentation is saved to the memory cell and the regressed level is also saved. Then the next vertebra segmentation is performed, and the centering process is repeated. The algorithm terminates when one of a several different conditions are met. For instance, in top-down mode, the algorithm terminates if it reaches the bottom of the scan, if the predicted level is greater than 24, or if no new bone is found.

At the end of this process, the segmentation with the level closest to L1 is saved as the predicted level. To output the ordering for the final output segmentation labelmap we perform a max-likelihood calculation to obtain the most likely ordering, where the regression outputs are re-interpreted as likelihoods. A vector of 24 likelihood values is determined for each vertebra. For each segmented vertebra i we calculate its likelihood vector l^i using an unnormalized Gaussian function :

$$l_{j}^{i} = e^{-\frac{1}{2}((L_{\text{pred}} - j)/\sigma)^{2}}$$
(2)

We set $\sigma = 2$ (the value used here does not seem to be critical). Next, all possible orderings consistent with the number of vertebra found are considered. If we store the ordering in a vector v of length N, then the likelihood for that ordering is computed as $\mathcal{L}(v) = \prod_{i=1}^{N} l_{v_i}^i$. The ordering with the highest likelihood is chosen as the final ordering.

3. RESULTS

3.1 Identifying the L1 level

Results for different tests are shown in table 2. The test set was a set of 40 scans from a completely different dataset of CT colonography (CTC) cases from the University of Wisoconsin Medical Center. When training with 44 hand segmentations the highest accuracy we were able to achieve was 70%. In roughly 15% of cases the L1 vertebra would be shifted up to T12 and in 5% it would be shifted down to L2 (additionally, around 5% of cases failed, often due to implants or other abnormalities). We then tried doing a 2 class segmentation of L1

Table 2. Results from different tests. For the 3D U-Net, The size of the validation sets varied between N = 4 to N = 10. Naug refers to the number of augmented versions generated for each scan. The average error and median error for the L1 level in the test set are reported in mm. Note this is not a head-to-head comparison as the dataset used for training the 3D U-Net was a different dataset using pseudosegmentations compared to the VerSe data used for training the instance segmentation approach. * 62 scans were used to generate 8,595 training volumes, including all augmented versions.

method	classes	N_{train}	N_{aug}	DICE _{val}	$avg err_{test} (mm)$	med err $_{\rm test}~({\rm mm})$	L1 acc.test
3D U-Net	(L1)	22	256	0.60	36.9	24.5	40%
	(L1, T12, ribs)	22	256	0.62	25.0	12.5	65%
	(L1, T12 + ribs)	22	256	0.70	22.0	15.9	73%
	(L1, T12)	28	256	0.50	24.4	18.1	59%
	(L1)	30	256	0.80	25.8	12.0	60%
	(L1)	38	256	0.80	16.4	9.0	65%
	(L1)	44	256	0.80	14.3	5.8	70%
	(L1)	189	8	0.70	8.0	3.0	88%
	(L1)	215	8	0.88	21.0	5.7	77%
	(L1)	249	8	0.71	11.2	4.5	98%
Instance	1 generic class	62^{*}	n/a	0.85	65.3	35.5	30%

and T12 and 2 class segmentations of L1 and T12+rib. However, the multiclass segmentations showed the same failure mode of "shifting up" as the single class, so we abandoned that approach.

We also experimented with a 6 layer FCNN to regress the distance to the L1 level from a subvolume but found that the performance of the network in the validation data during training plateaued with an unacceptably high error of 30 mm.



Figure 2. Segmentations of L1 from the U-Net trained on 189 pseudo-segmentations, showing the axial, sagittal, and coronal views. The left image is a CTC case, while the right one is a hold-out case from the Naval data, the dataset used for training.

3.2 Iterative instance segmentation

We first tried to train the iterative instance based approach on the same type of pseudosegmentations we used to train the L1 segmentation U-Net. The algorithm requires that the output threshold be tuned afterwards for optimum performance. Several segmentations, which are randomly drawn (ie. not hand picked) are shown in figure 3. One issue is apparently is that the method often has trouble segmenting the first vertebra because the instance memory is empty and cases with an empty instance memory are rare in the training dataset. One possible trick to overcome this problem would be to to train two separate models for both bottom-up and one for top-down inference, and then combine the bottom-up and top-down segmentation outputs. Another solution would be to create additional training data with the vertebra to be segmented at the top or bottom of the subvolume and an empty instance memory. Quantitative results on the test set of 40 CTC cases are presented in table 2. While this method is not as accurate at identifying L1, qualitatively the segmentations in the CTC and other test datasets are of higher quality than with the direct approach. This is not surprising given that the instance based U-Net was exposed to many more examples of vertebra during training.



Figure 3. Segmentations of the spine produced using the iterative-instance based method randomly chosen from four datasets. From left to right: VerSe sample (in-sample ground truth dataset), CT Colonography case from the University of Wisconsin Medical Center, CT Colonography case from National Navy Medical Center,²² and a renal donor case in the post-contrast phase. While the algorithm shows good generalization performance to the different dataset and is capable of working on scans with arbitrary field of view, several weaknesses are apparent such as trouble segmenting the topmost and bottomost vertebrae.

4. CONCLUSION

We have compared two very different approaches to identifying the L1 level. First we trained a 3D U-Net to segment the L1 vertebra directly. As far as we know this is the first time this approach has been studied. By training with 249 pseudosegmentations we obtained a model which achieved an accuracy of 98% at identifying L1 in an independent test dataset of 40 scans. We also developed an iterative instance based segmentation system to segment and classify the entire spine. We hypothesized this approach may be useful when there are vertebral anomalies present in the scan such as additional or missing vertebra. We first trained using the same pseudosegmentations used for training the U-Net for direct segmentation, but we found the results to be of poor quality. Next we trained with hand drawn segmentations from the VerSe dataset. The resulting model yielded more accurate segmentations of nearly the entire spine but a less accurate classification accuracy for L1. While our main goal here was to identify L1, the high quality segmentations obtained by the iterative instance based approach may be useful for performing a measurement of bone mineral density, for instance by performing an erosion on the segmentation and averaging the Hounsfield units within the trabecular space of the L1 vertebra, similar to Liu et al.³¹

ACKNOWLEDGMENTS

We thank Dr. Perry J. Pickhardt and Dr. J. Richard Choi for providing the CTC images used for testing. This research was supported in part by the Intramural Research Program of the National Institutes of Health Clinical Center.

REFERENCES

- [1] Burns, J. E., Yao, J., and Summers, R. M., "Artificial intelligence in musculoskeletal imaging: A paradigm shift," *Journal of Bone and Mineral Research* (Aug. 2019).
- [2] Sahiner, B., Pezeshk, A., Hadjiiski, L. M., Wang, X., Drukker, K., Cha, K. H., Summers, R. M., and Giger, M. L., "Deep learning in medical imaging and radiation therapy," *Medical Physics* 46, e1–e36 (Nov. 2018).
- [3] Jang, S., Graffy, P. M., Ziemlewicz, T. J., Lee, S. J., Summers, R. M., and Pickhardt, P. J., "Opportunistic osteoporosis screening at routine abdominal and thoracic CT: Normative L1 trabecular attenuation values in more than 20 000 adults," *Radiology* 291, 360–367 (May 2019).
- [4] Pickhardt, P. J., Lee, S. J., Liu, J., Yao, J., Lay, N., Graffy, P. M., and Summers, R. M., "Population-based opportunistic osteoporosis screening: Validation of a fully automated CT tool for assessing longitudinal BMD changes," *The British Journal of Radiology* 92, 20180726 (Feb. 2019).

- [5] Lee, S. J., Liu, J., Yao, J., Kanarek, A., Summers, R. M., and Pickhardt, P. J., "Fully automated segmentation and quantification of visceral and subcutaneous fat at abdominal CT: application to a longitudinal adult screening cohort," *The British Journal of Radiology*, 20170968 (Mar. 2018).
- [6] Graffy, P. M., Liu, J., Pickhardt, P. J., Burns, J. E., Yao, J., and Summers, R. M., "Deep learning-based muscle segmentation and quantification at abdominal CT: application to a longitudinal adult screening cohort for sarcopenia assessment," *The British Journal of Radiology* 92, 20190327 (Aug. 2019).
- [7] Roth, H. R., Lu, L., Lay, N., Harrison, A. P., Farag, A., Sohn, A., and Summers, R. M., "Spatial aggregation of holistically-nested convolutional neural networks for automated pancreas localization and segmentation," *Medical Image Analysis* 45, 94–107 (Apr. 2018).
- [8] Graffy, P. M., Liu, J., O'Connor, S., Summers, R. M., and Pickhardt, P. J., "Automated segmentation and quantification of a ortic calcification at abdominal CT: application of a deep learning-based algorithm to a longitudinal screening cohort," *Abdominal Radiology* 44, 2921–2928 (Apr. 2019).
- [9] Weston, A. D., Korfiatis, P., Kline, T. L., Philbrick, K. A., Kostandy, P., Sakinis, T., Sugimoto, M., Takahashi, N., and Erickson, B. J., "Automated abdominal segmentation of CT scans for body composition analysis using deep learning," *Radiology* 290, 669–679 (Mar. 2019).
- [10] French, H. D., Somasundaram, A. J., Schaefer, N. R., and Laherty, R. W., "Lumbosacral transitional vertebrae and its prevalence in the australian population," *Global Spine Journal* 4, 229–232 (Sept. 2014).
- [11] Aly, I., Chapman, J. R., Oskouian, R. J., Loukas, M., and Tubbs, R. S., "Lumbar ribs: a comprehensive review," *Childs Nervous System* **32**, 781–785 (Sept. 2015).
- [12] Glass, R. B. J., Norton, K. I., Mitre, S. A., and Kang, E., "Pediatric ribs: A spectrum of abnormalities," *RadioGraphics* 22, 87–104 (Jan. 2002).
- [13] Lessmann, N., van Ginneken, B., de Jong, P. A., and Išgum, I., "Iterative fully convolutional neural networks for automatic vertebra segmentation and identification," *Medical Image Analysis* 53, 142–155 (Apr. 2019).
- [14] Išgum, I., van Ginneken, B., and Lessmann, N., "Iterative convolutional neural networks for automatic vertebra identification and segmentation in CT images," in [Medical Imaging 2018: Image Processing], Angelini, E. D. and Landman, B. A., eds., SPIE (Mar. 2018).
- [15] McCouat, J. and Glocker, B., "Vertebrae detection and localization in ct with two-stage cnns and dense annotations," *arXiv preprint arXiv:1910.05911*.
- [16] Liao, H., Mesfin, A., and Luo, J., "Joint vertebrae identification and localization in spinal CT images by combining short- and long-range contextual information," *IEEE Transactions on Medical Imaging* 37, 1266–1275 (May 2018).
- [17] Yang, D., Xiong, T., Xu, D., Huang, Q., Liu, D., Zhou, S. K., Xu, Z., Park, J., Chen, M., Tran, T. D., Chin, S. P., Metaxas, D., and Comaniciu, D., "Automatic vertebra labeling in large-scale 3d CT using deep image-to-image network with message passing and sparsity regularization," in [Lecture Notes in Computer Science], 633–644, Springer International Publishing (2017).
- [18] Janssens, R., Zeng, G., and Zheng, G., "Fully automatic segmentation of lumbar vertebrae from CT images using cascaded 3d fully convolutional networks," in [2018 IEEE 15th International Symposium on Biomedical Imaging (ISBI 2018)], IEEE (Apr. 2018).
- [19] Roth, H. R., Lu, L., Seff, A., Cherry, K. M., Hoffman, J., Wang, S., Liu, J., Turkbey, E., and Summers, R. M., "CT Lymph Nodes Dataset- The Cancer Imaging Archive." http://doi.org/10.7937/K9/TCIA. 2015.AQIIDCNM (2015).
- [20] Roth, H. R., Lu, L., Seff, A., Cherry, K. M., Hoffman, J., Wang, S., Liu, J., Turkbey, E., and Summers, R. M., "A new 2.5D representation for lymph node detection using random sets of deep convolutional neural network observations," in [Medical Image Computing and Computer-Assisted Intervention – MICCAI 2014], 520–527, Springer International Publishing (2014).
- [21] Clark, K., Vendt, B., Smith, K., Freymann, J., Kirby, J., Koppel, P., Moore, S., Phillips, S., Maffitt, D., Pringle, M., Tarbox, L., and Prior, F., "The cancer imaging archive (TCIA): Maintaining and operating a public information repository," *Journal of Digital Imaging* 26, 1045–1057 (July 2013).
- [22] Pickhardt, P. J., Choi, J. R., Hwang, I., Butler, J. A., Puckett, M. L., Hildebrandt, H. A., Wong, R. K., Nugent, P. A., Mysliwiec, P. A., and Schindler, W. R., "Computed tomographic virtual colonoscopy to screen for colorectal neoplasia in asymptomatic adults," *New England Journal of Medicine* **349**, 2191–2200 (Dec. 2003).

- [23] Yao, J., O'Connor, S., and Summers, R., "Automated spinal column extraction and partitioning," in [3rd IEEE International Symposium on Biomedical Imaging: Macro to Nano, 2006.], IEEE.
- [24] Sekuboyina, A., Rempfler, M., Kukačka, J., Tetteh, G., Valentinitsch, A., Kirschke, J. S., and Menze, B. H., "Btrfly net: Vertebrae labelling with energy-based adversarial learning of local spine prior," in [Medical Image Computing and Computer Assisted Intervention – MICCAI 2018], 649–657, Springer International Publishing (2018).
- [25] Ibragimov, B., Likar, B., Pernus, F., and Vrtovec, T., "Shape representation for efficient landmark-based segmentation in 3-d," *IEEE Transactions on Medical Imaging* 33, 861–874 (Apr. 2014).
- [26] Korez, R., Ibragimov, B., Likar, B., Pernus, F., and Vrtovec, T., "A framework for automated spine and vertebrae interpolation-based detection and model-based segmentation," *IEEE Transactions on Medical Imaging* 34, 1649–1662 (Aug. 2015).
- [27] Ronneberger, O., Fischer, P., and Brox, T., "U-net: Convolutional networks for biomedical image segmentation," in [Lecture Notes in Computer Science], 234–241, Springer International Publishing (2015).
- [28] Ozgün Çiçek, Abdulkadir, A., Lienkamp, S. S., Brox, T., and Ronneberger, O., "3D U-Net: Learning dense volumetric segmentation from sparse annotation," in [Medical Image Computing and Computer-Assisted Intervention – MICCAI 2016], 424–432, Springer International Publishing (2016).
- [29] Sandfort, V., Yan, K., Pickhardt, P. J., and Summers, R. M., "Data augmentation using generative adversarial networks (CycleGAN) to improve generalizability in CT segmentation tasks," *Scientific Reports* 9 (Nov. 2019).
- [30] Wu, Y. and He, K., "Group normalization," in [Computer Vision ECCV 2018], 3–19, Springer International Publishing (2018).
- [31] Liu, S., Gonzalez, J., Zulueta, J., de Torres, J. P., Yankelevitz, D. F., Henschke, C. I., and Reeves, A. P., "Fully automated bone mineral density assessment from low-dose chest CT," in [Medical Imaging 2018: Computer-Aided Diagnosis], Mori, K. and Petrick, N., eds., SPIE (Feb. 2018).