

RibSeg Dataset and Strong Point Cloud Baselines for Rib Segmentation from CT Scans

Jiancheng Yang^{1,2,*}, Shixuan Gu^{1,*}, Donglai Wei³,
Hanspeter Pfister³, and Bingbing Ni^{1†}

¹ Shanghai Jiao Tong University, Shanghai, China
nibingbing@sjtu.edu.cn

² Dianei Technology, Shanghai, China

³ Harvard University, Cambridge MA, USA

Abstract. Manual rib inspections in computed tomography (CT) scans are clinically critical but labor-intensive, as 24 ribs are typically elongated and oblique in 3D volumes. Automatic rib segmentation methods can speed up the process through rib measurement and visualization. However, prior arts mostly use in-house labeled datasets that are publicly unavailable and work on dense 3D volumes that are computationally inefficient. To address these issues, we develop a labeled rib segmentation benchmark, named *RibSeg*, including 490 CT scans (11,719 individual ribs) from a public dataset. For ground truth generation, we used existing morphology-based algorithms and manually refined its results. Then, considering the sparsity of ribs in 3D volumes, we thresholded and sampled sparse voxels from the input and designed a point cloud-based baseline method for rib segmentation. The proposed method achieves state-of-the-art segmentation performance (Dice \approx 95%) with significant efficiency (10 \sim 40 \times faster than prior arts). The RibSeg dataset, code, and model in PyTorch are available at <https://github.com/M3DV/RibSeg>.

Keywords: rib segmentation · rib centerline · medical image dataset · point clouds · computed tomography.

1 Introduction

Detection and diagnosis of rib-related diseases, *e.g.*, rib fracture and bone lesions, is essential and common in clinical practice and forensics. For instance, rib fracture detection can identify chest trauma severity that accounts for 10% \sim 15% of all traumatic injuries [17], and bone metastases are common in solid tumours [3]. Chest computed tomography (CT) is a primary choice for examining chest trauma thanks to its advantage in revealing occult fractures and fracture-related complications [6]. However, understanding ribs in CT imaging is challenging: 24 ribs in human bodies are typically elongated and oblique in 3D volumes, with complex geometries across numerous CT sections; in other words, a large number

*These authors have contributed equally: Jiancheng Yang and Shixuan Gu.

†Corresponding author: Bingbing Ni (nibingbing@sjtu.edu.cn).




Fig. 1: **Illustration of RibSeg Dataset:** 490 CT scans from the existing RibFrac dataset [7] and new annotations for rib segmentation. Colors in rib segmentation figures denotes the axial depth.

of CT slices must be evaluated sequentially, rib-by-rib and side-by-side, which is tedious and labor-intensive for clinicians. Besides, fractures/lesions could be inconspicuous when reading 2D slices. These challenges urge the development and application of rib visualization tools, *e.g.*, rib unfolding [15], whose core technique is (automatic) rib segmentation and centerline extraction.

A few prior arts have addressed rib segmentation [16,18,8] before the era of deep learning. Rib tracing is a popular method, while it is highly sensitive to initially detected seed points and vulnerable to local ambiguities. Supervised deep learning-based segmentation [9] from CT volumes is robust as it learns hierarchical visual features from raw voxels; However, this study does not consider the sparsity and elongated geometry of ribs in 3D volumes. There are also studies focusing on rib centerline extraction instead of full rib segmentation, *e.g.*, rib tracing [14] and deformable template matching [21]. This study focuses on rib segmentation, where rib centerlines could be extracted with geometric post-processing algorithms.

Although researchers have made progress in rib segmentation and centerline extraction, there is no public dataset in this topic, making it difficult to benchmark existing methods and develop downstream applications (*e.g.*, rib fracture detection). Besides, existing methods work on the dense 3D volumes instead of the sparse rib voxels, which are computationally inefficient: around 5s to 20s to segment ribs [9] or 40s to extract rib centerlines [21]. To address these issues, we first develop a large-scale CT dataset, named *RibSeg* for rib segmentation (Fig. 1). The CT scans come from the public RibFrac dataset [7] consisting of 660 chest-abdomen CT scans for rib fracture segmentation, detection, and classification. As ribs are relatively recognizable compared to other anatomical structures, we use hand-crafted 3D image processing algorithms to generate segmentation and then manually check and refine the labels. This procedure, with little annotation effort, generates visually satisfactory rib segmentation labels for 490 CT scans. Moreover, considering the sparsity of ribs in 3D volumes ($< 0.5\%$ voxels) and high HU values (> 200) of bone structures in CT scans, we propose an efficient point cloud-based model to segment ribs, which is an order of magnitude faster than previous methods. The proposed method converts dense CT volume into sparse point clouds via thresholding and random down-sampling, and produces high-quality and robust rib segmentation (Dice $\approx 95\%$).

We can further extract rib centerlines from the predicted rib segmentation with post-processing.

The RibSeg dataset could be used for the development of downstream rib-related applications. Besides, considering the differences from standard medical image datasets [1,23] with pixel/voxel grids, the elongated shapes and oblique poses of ribs enable the RibSeg dataset to serve as a benchmark for curvilinear structures and geometric deep learning (*e.g.*, point clouds).

Contributions. 1) The first public benchmark for rib segmentation, which enables downstream applications and method comparison. 2) A novel point-based perspective on modeling 3D medical images beyond voxel grids. 3) A point cloud-based rib segmentation baseline with high efficiency and accuracy.

2 Materials and Methods

2.1 RibSeg Dataset

Dataset Overview. The RibSeg dataset uses the public computed tomography (CT) scans from RibFrac dataset [7], an open dataset with 660 chest-abdomen CT scans for rib fracture segmentation, detection, and classification. The CT scans are saved in NIFTI (.nii) format with volume sizes of $N \times 512 \times 512$, where 512×512 is the size of CT slices, and N is the number of CT slices (typically $300 \sim 400$). Most cases are confirmed with complete rib cages and manually annotated with at least one rib fracture by radiologists.

As ribs are relatively recognizable compared to other anatomical structures, we use a semi-automatic approach (see details in the following section) to generate rib segmentation, with hand-crafted morphology-based image processing algorithms, as well as manual checking and refinement. Though computationally intensive, this approach produces visually satisfactory labels with few annotation efforts. Finally, there are 490 qualified CT cases in the RibSeg dataset with 11,648 individual ribs in total, where each case has segmentation labels of 24 (or 22 in some cases) ribs. We also provide the rib centerline ground truth extracted from the rib segmentation labels. Note that the rib segmentation and centerline ground truth are imperfect, as the annotations are generated with algorithms. Besides, only voxels with higher HU than a threshold (200) are included, making the rib segmentation annotations in the RibSeg dataset hollow. However, we manually check the ground truth labels for rib segmentation and centerlines to ensure that the included 490 datasets are high-quality enough to develop downstream applications. The data split of the RibSeg dataset is summarized in Table 1: training set (320 cases, to train the deep learning system), development set (a.k.a validation set, 50 cases, to tune hyperparameters of the deep learning system), and test set (120 cases, to evaluate the model). The RibSeg training, development, and test set are from those of the RibFrac dataset respectively, enabling the development of downstream applications (*e.g.*, rib fracture detection) in the MICCAI 2020 RibFrac challenge¹.

¹<https://ribfrac.grand-challenge.org/>

Table 1: Overview of RibSeg dataest.

Subset	No. of CT Scans	No. of Individual Ribs
Training / Development / Test	320 / 50 / 120	7,670 / 1,187 / 2,862

Semi-Automatic Annotation and Quality Control. We describe the primary steps of the annotation procedure as follows:

Rib Segmentation. For each volume, we first filter out non-target voxels by thresholding and removing regions outside the bodies. Considering the geometric differences between ribs and vertebra, we separate the ribs from vertebra using morphology-based image processing algorithms (*e.g.*, dilation, erosion). In some cases, the segmentation result contains parts of the clavicle and scapula. Therefore, we manually locate those non-target voxels and remove them according to the coordinates of their connected components.

Centerline Extraction. Based on the rib segmentation, we extract the centerline by implementing the following procedure on each rib (connected component): randomly select two points at both ends of the cylinder dilated from the rib, calculate the shortest path between the points¹ [19], and smoothen the path to obtain centerline. This procedure produces high-quality centerlines even from coarse rib segmentation. At the end of extraction, we label both centerlines and rib segmentation in the order of top to bottom and left to right.

Manual Checking and Refinement. The abnormal cases, along with the pursuit of high annotation quality, motivate us to perform laborious checking and refinement after both rib segmentation and centerline extraction stages. For instance, a few cases miss floating ribs after segmentation, which reduces the connected components in annotation to 22 or less. Hence we have to check and refine the annotation case by case manually. To recover and annotate missed ribs, we turn back to the previous stage to ensure segmentation completeness by modifying the corresponding components.

2.2 Rib Segmentation from a Viewpoint of Point Clouds

The key insight of the proposed method is that simple algorithm (*i.e.*, thresholding in this study) can produce the candidate voxels for bone structures. Thus, we can avoid heavy computation on dense voxels with sparse point clouds instead. Besides, the point cloud methods use geometric information directly, reducing the texture bias of pixel/voxel-based CNNs [4]. The point cloud viewpoint has the potential to generalize to other anatomical structures whose coarse prediction could be obtained cheaply.

Deep Learning for Point Cloud Analysis. Deep learning for point cloud analysis [5] has been an emerging field in computer vision thanks to the popularity of 3D sensors (*e.g.*, LiDAR and RGB-D cameras) and computer graphics

¹<https://github.com/pangyuteng/simple-centerline-extraction>




Fig. 2: **Rib Segmentation from a Viewpoint of Point Clouds.** The CT volumes are first binarized to obtain candidate bone voxels as inputs, then a point cloud neural network (e.g., PointNet++ [13]) is used to segment each point in downsampled input point clouds. Note that the downsampling scale is different during training (30K points) and inference (250K points).

applications (e.g., games, VR/AR/MR), PointNet [12] and DeepSet [26] pioneer this direction, where a symmetric function (e.g., shared FC) is used for learning high-level representation before aggregation (e.g., pooling). Following studies introduces sophisticated feature aggregation based on spatial graphs [13,10] or attention [24]. However, only a few studies have applied deep shape analysis in medical imaging scenarios [25,20,22].

Model Pipeline. Considering the sparsity of ribs in 3D volumes ($< 0.5\%$ voxels) and high HU values of bones in CT scans, we design a point cloud-based model to segment ribs on binarized sparse voxels. As depicted in Fig. 2, we first set a threshold of 200 HU (Hounsfield Unit) to filter out the non-bone voxels roughly. The resultant binarized volumes are randomly downsampled and converted to point sets for ease of computation before forwarding to the network.

Our point cloud-based model is expected to infer dense predictions from large-scale point sets, which has to address the memory issue. Hence a custom PointNet++ [13], with its adjustable memory footprint, is adopted as backbone. Capable of learning local features with increasing contextual scales, PointNet++ has shown compelling robustness on sparse 3D point cloud segmentation tasks. Through set abstraction, geometric features of ribs can be extracted from binarized sparse voxels facilitating the rib segmentation task. For post-processing, the model output point prediction is converted back to volumes (voxel prediction) by morphology-based image processing algorithms.

Model Training and Inference. During the training stage, batches are downsampled to 30,000 points per volume considering the trade-off between batch size and input size. We apply online data augmentations, including scaling, translation, and jittering, to all downsampled point sets before forwarding them to the neural network. The Adam optimizer is adopted to train all models end-to-end for 250 epochs with the batch size of 8 and cross-entropy loss (CE) as the loss function. The initial learning rate was set at 0.001 and decayed by a factor of 0.5 every 20 epochs with the lower bound of 10^{-5} .

During the inference stage, volumes are converted to point sets with the size of 250,000. The model then produces dense point predictions on rib segmenta-

tion. The point predictions are converted back to dense volumes by dilation, and we obtain the voxel predictions of rib segmentation by taking the intersections between the dilated volumes and the binarized volumes.

Model Evaluation. As the model only outputs sparse point predictions, we post-process the point predictions back to volumes (dense voxel predictions). Both point-wise and voxel-wise segmentation performance is evaluated,

$$Dice^{(L)} = 2 \cdot |y_{(L)} \cdot \hat{y}_{(L)}| / (|y_{(L)}| + |\hat{y}_{(L)}|), \quad L \in \{P, V\}, \quad (1)$$

where $Dice^{(P,V)}$ indicates the sparse point-wise and dense voxel-wise Dice.

Apart from segmentation performance, we also report the missing ratio of individual ribs to evaluate the clinical applicability. Specifically, a missing of an individual rib i is counted if $recall_i < 0.5$, and then the missing ratio can be calculated with ease. As the segmentation of first and twelfth rib pairs tend to be more difficult, we calculate and report the missing ratio of all/first/intermediate/twelfth rib pairs, as depicted in Table 2.

3 Results

3.1 Quantitative Analysis

For model accuracy comparison, we first implement a 3D UNet [2] taking patches of CT volumes as input with the same setting of FracNet [7]. Moreover, we train two models with and without data augmentation, respectively. The models are evaluated with two input (point sets) sizes: 30K (input size in the training stage) and 250K. As point-wise Dice is only a proxy metric, we focus on voxel-wise Dice, as it is fair for any methods in rib segmentation. As depicted in Table 2, all point-based methods significantly outperform voxel-based 3D UNet; Besides, as point-based methods take whole volumes as inputs, it is more efficient than voxel-based method. The methods with data augmentation are at least 2% higher than the methods without data augmentation, and methods with large-scale input enjoy 0.9%~1.3% higher values, as dense point prediction leads to rich details in voxel prediction. When it comes to the missing ratio of ribs, the method with data augmentation and a large input size performs best. The comparison results show that training-time data augmentation and inference-time large input volume size can improve the result. These quantitative metrics also indicate the potentials of our method in clinical applications.

In terms of the run-time, point-based methods have a clear advantage. While methods with a 250K-point input size have a little bit higher time consumption (0.8s), it is acceptable in consideration of its performance boost. The inference time was measured with the implementation of PyTorch 1.7.1 [11] and Python 3.7, on a machine with a single NVIDIA Tesla P100 GPU with Intel(R) Xeon(R) CPU @ 2.20 GHz and 150 G memory. As a reference, prior art [9] based on 3D networks takes a model forward time of 5~20 seconds, with a Dice value of 84% at best. The previous work [21] takes 40 seconds to extract the rib centerline.

Table 2: Quantitative metrics on RibSeg test set, including Dice over sparse points ($Dice^{(P)}$), Dice over dense voxels after post-processing ($Dice^{(V)}$), ratio of missed all/first/intermediate/twelfth rib pairs (A/F/I/T) at recall > 0.5, and the model forward time in second. Post-processing time is not included as it heavily depends on the implementation.

Methods	$Dice^{(P)}$	$Dice^{(V)}$	Missed Ribs (A/F/I/T)	Forward (s)
Voxel-Based 3D UNet [2,7]	-	86.3%	4.6%/7.9%/2.3%/24.6%	30.63
PN++ [13] (30K)	92.3%	91.0%	1.6%/2.9%/0.7%/10.4%	0.32
PN++ [13] (250K)	91.5%	92.3%	0.9%/3.3%/0.3%/4.7%	1.12
PN++ [13] (30K) + aug.	94.9%	94.3%	1.1%/0.8%/0.4%/9.0%	0.32
PN++ [13] (250K) + aug.	94.6%	95.2%	0.6%/0.4%/0.2%/5.2%	1.12

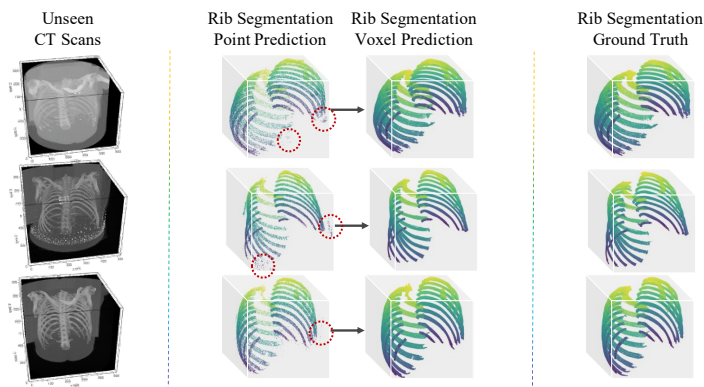


Fig. 3: **Visualization of Predicted Rib Segmentation.** Red circles denote imperfect (sparse) point prediction.

However, a direct comparison of metrics and speed is unfair since these results were measured with different infrastructures on different datasets. Note that, post-processing time is not included as it heavily depends on the implementation (*e.g.*, programming languages, parallel computing).

3.2 Qualitative Analysis

Visualization on Predicted Rib Segmentation. Fig. 3 visualizes the point-level and voxel-level rib segmentation prediction. As depicted in Fig. 3, point predictions of the first 3 CT scans are visually acceptable, which can smoothly produce voxel-level segmentation predictions. The results are promising, and it is even hard to tell the visual difference between predictions and ground truths.

As depicted in Fig. 3, the point prediction is imperfect even with high segmentation Dice. After manually-tuned post-processing, these imperfect predictions could be fixed to some degree. For instance, the point prediction of the second

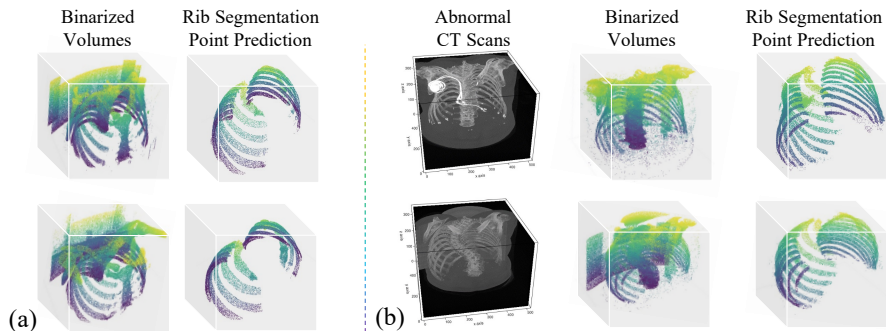


Fig. 4: **Robustness Test on Extreme Cases.** (a) Point prediction on unseen incomplete rib cages. (b) Point prediction on unseen abnormal CT scans.

CT scan contains a certain part of the scapula that can be nicely filtered during post-processing. However, the point prediction of the third CT scan suffers a missing on the first pair of ribs, which can not be ignored. The post-process is not able to fix it and produces incomplete voxel segmentation. Despite the small missing part, the rib cages in the predictions on rib segmentation are still visually acceptable.

Robustness Test on Extreme Cases. To evaluate the robustness of the proposed point cloud-based model, three kinds of extreme cases are selected for inference: CT scans of incomplete rib cages, CT scans of serious spinal pathology, and CT scans containing metal objects inseparable from ribs (*e.g.*, pacemaker). Incomplete CT scans are rather common in clinical cases, which makes the tests practically critical. As depicted in Fig. 4 (a), we randomly select unseen CT scans and take the upper half of their binarized volumes for inference. Delightfully, the point predictions are visually qualified for clinical applications.

For further robustness evaluation, we test on a case of serious spinal pathology and a case containing a pacemaker of high density. Regarding the case with a pacemaker, it is inseparable from ribs, which makes the segmentation extremely laborious to obtain. Hence we save the trouble of manual segmentation by setting it as an abnormal case when building the RibSeg Dataset. As depicted in Fig. 4 (b), our prediction on the case of spinal pathology looks complete as if it is nicely segmented. While the prediction on the other case contains a small number of voxels belong to the pacemaker, the whole rib cage is well segmented. Considering that all cases selected for the robustness test are unseen and geometrically difficult to segment, such promising prediction results may confirm the strong robustness of the proposed method as well as its potentials to be clinically applicable.

Post-Processing Rib Centerlines. With the high-quality rib segmentation, as illustrated in Fig. 5, the rib centerlines could be obtained by post-processing




Fig. 5: **Post-Processing Rib Centerlines** from predicted rib segmentation.

(*i.e.*, shortest path between points, same as the procedure with ground truth in Sec. 2.1). Although not end-to-end, the rib centerline predictions are visually acceptable for most cases. However, the post-processing algorithms for rib centerlines are sensitive to rib fractures and other abnormal cases. Considering the high clinical importance of rib centerlines, it urges a more robust method for rib centerline extraction with rib segmentation and centerline labels provided by the RibSeg dataset.

4 Conclusion and Further Work

We built the RibSeg dataset, which is the first open dataset for rib segmentation. On this dataset, we benchmarked a point cloud-based method with high performance and significant efficiency. The proposed method shows potentials to be clinically applicable, enhancing the efficiency and performance of downstream tasks, such as the diagnosis of rib fractures and bone lesions. Besides the clinical application, the RibSeg dataset could also serve as an interesting benchmark for curvilinear structures and geometric deep learning (*e.g.*, point clouds), considering the special geometry of rib structures.

There are several limitations in this study. The annotations in this paper are generated with hand-crafted morphological algorithms, and then manually checked by a junior radiologist with 3D Slicers. While such pipeline reduces the annotation cost, it cannot handle cases when the initial automatic method fails. Thus, we only managed to annotate the segmentation for 490 cases out of the 660 cases in the RibFrac dataset. Also, for the centerline extraction task that is essential for rib-related applications, we take a two-stage approach and apply heavy post-processing method to the first-stage segmentation result. Such approach is sensitive to rib fractures and segmentation errors for other abnormal cases and a more robust method will be favorable.

Acknowledgment. This work was supported by the National Science Foundation of China (U20B2072, 61976137).

References

1. Antonelli, M., Reinke, A., Bakas, S., et al.: The medical segmentation decathlon. arXiv preprint arXiv:2106.05735 (2021)

2. Çiçek, Ö., Abdulkadir, A., Lienkamp, S.S., Brox, T., Ronneberger, O.: 3d u-net: learning dense volumetric segmentation from sparse annotation. In: MICCAI. pp. 424–432. Springer (2016)
3. Coleman, R., Body, J.J., Aapro, M., Hadji, P., Herrstedt, J., Group, E.G.W., et al.: Bone health in cancer patients: Esmo clinical practice guidelines. *Annals of oncology* **25**, iii124–iii137 (2014)
4. Geirhos, R., Rubisch, P., Michaelis, C., Bethge, M., Wichmann, F.A., Brendel, W.: Imagenet-trained cnns are biased towards texture; increasing shape bias improves accuracy and robustness. In: ICLR (2019)
5. Guo, Y., Wang, H., Hu, Q., Liu, H., Liu, L., Bennamoun, M.: Deep learning for 3d point clouds: A survey. *IEEE transactions on pattern analysis and machine intelligence* (2020)
6. Jin, L., Ge, X., Lu, F., Sun, Y., et al.: Low-dose ct examination for rib fracture evaluation: A pilot study. *Medicine* **97**(30) (2018)
7. Jin, L., Yang, J., Kuang, K., Ni, B., Gao, Y., Sun, Y., Gao, P., Ma, W., Tan, M., Kang, H., et al.: Deep-learning-assisted detection and segmentation of rib fractures from ct scans: Development and validation of fracnet. *EBioMedicine* **62**, 103106 (2020)
8. Klinder, T., Lorenz, C., Von Berg, J., Dries, S.P., Bülow, T., Ostermann, J.: Automated model-based rib cage segmentation and labeling in ct images. In: MICCAI. pp. 195–202. Springer (2007)
9. Lenga, M., Klinder, T., Bürger, C., von Berg, J., Franz, A., Lorenz, C.: Deep learning based rib centerline extraction and labeling. In: MICCAI International Workshop on Computational Methods and Clinical Applications in Musculoskeletal Imaging. pp. 99–113. Springer (2018)
10. Liu, J., Ni, B., Li, C., Yang, J., Tian, Q.: Dynamic points agglomeration for hierarchical point sets learning. In: ICCV. pp. 7546–7555 (2019)
11. Paszke, A., Gross, S., Chintala, S., Chanan, G., Yang, E., Devito, Z., Lin, Z., Desmaison, A., Antiga, L., Lerer, A.: Automatic differentiation in pytorch (2017)
12. Qi, C.R., Su, H., Mo, K., Guibas, L.J.: Pointnet: Deep learning on point sets for 3d classification and segmentation. In: CVPR. pp. 652–660 (2017)
13. Qi, C.R., Yi, L., Su, H., Guibas, L.J.: Pointnet++: Deep hierarchical feature learning on point sets in a metric space. In: NIPS (2017)
14. Ramakrishnan, S., Alvino, C., Grady, L., Kiraly, A.: Automatic three-dimensional rib centerline extraction from ct scans for enhanced visualization and anatomical context. In: *Medical Imaging 2011: Image Processing*. vol. 7962, p. 79622X (2011)
15. Ringl, H., Lazar, M., Töpker, M., Woitek, R., Prosch, H., Asenbaum, U., Balassy, C., Toth, D., Weber, M., Hajdu, S., et al.: The ribs unfolded—a ct visualization algorithm for fast detection of rib fractures: effect on sensitivity and specificity in trauma patients. *European radiology* **25**(7), 1865–1874 (2015)
16. Shen, H., Liang, L., Shao, M., Qing, S.: Tracing based segmentation for the labeling of individual rib structures in chest ct volume data. In: MICCAI. pp. 967–974. Springer (2004)
17. Sirmali, M., et al.: A comprehensive analysis of traumatic rib fractures: morbidity, mortality and management. *European Journal of Cardio-Thoracic Surgery* **24**(1), 133–138 (2003)
18. Staal, J., van Ginneken, B., Viergever, M.A.: Automatic rib segmentation and labeling in computed tomography scans using a general framework for detection, recognition and segmentation of objects in volumetric data. *Medical image analysis* **11**(1), 35–46 (2007)

19. Teng, P.y., Bagci, A.M., Alperin, N.: Automated prescription of an optimal imaging plane for measurement of cerebral blood flow by phase contrast magnetic resonance imaging. *IEEE transactions on biomedical engineering* **58**(9), 2566–2573 (2011)
20. Wickramasinghe, U., Remelli, E., Knott, G., Fua, P.: Voxel2mesh: 3d mesh model generation from volumetric data. In: MICCAI. pp. 299–308. Springer (2020)
21. Wu, D., Liu, D., Puskas, Z., Lu, C., Wimmer, A., Tietjen, C., Soza, G., Zhou, S.K.: A learning based deformable template matching method for automatic rib centerline extraction and labeling in ct images. In: CVPR. pp. 980–987. IEEE (2012)
22. Yang, J., Fang, R., Ni, B., Li, Y., Xu, Y., Li, L.: Probabilistic radiomics: Ambiguous diagnosis with controllable shape analysis. In: MICCAI (2019)
23. Yang, J., Shi, R., Ni, B.: Medmnist classification decathlon: A lightweight automl benchmark for medical image analysis. In: ISBI (2021)
24. Yang, J., Zhang, Q., Ni, B., Li, L., Liu, J., Zhou, M., Tian, Q.: Modeling point clouds with self-attention and gumbel subset sampling. In: CVPR. pp. 3323–3332 (2019)
25. Yang, X., Xia, D., Kin, T., Igarashi, T.: Intra: 3d intracranial aneurysm dataset for deep learning. In: CVPR. pp. 2656–2666 (2020)
26. Zaheer, M., Kottur, S., Ravanbakhsh, S., Póczos, B., Salakhutdinov, R.R., Smola, A.J.: Deep sets. In: NIPS (2017)