Semi-supervised Learning with Contrastive and Topology Losses for Catheter Segmentation and Misplacement Prediction

Author name(s) withheld EMAIL(S) WITHHELD
Address withheld

Editors: Under Review for MIDL 2023

Abstract

Chest X-ray images are often used to determine the proper placement of catheters, as incorrect placement can lead to severe complications. With the advent of deep learning, computer-aided detection methods have been developed to assist radiologists in identifying catheter misplacement by detecting and highlighting the catheter's path. However, obtaining large, pixel-wise labeled datasets can be challenging due to the labor-intensive nature of annotation. To address this issue, we proposed a novel semi-supervised learning method that combines contrastive loss and topology loss. This method takes advantage of the known topological properties of catheters and does not require extensive labeling. We collected 7,378 chest X-ray images from the *****, which were labeled for misplacement of nasogastric and endotracheal tube catheters, and included pixel-wise annotation. Moreover, the CLiP dataset was used as an unlabeled dataset for semi-supervised learning. We used a hybrid U-Net architecture with an added classification head to perform simultaneous segmentation of the catheter and misplacement classification. Our model achieved an average AUC of 0.977 for classification and a average Dice score of 0.598 for segmentation.

Keywords: Semi-supervised Learning, Contrastive Learning, Topology Loss, Catheter Misplacement, Catheter Segmentation

1. Introduction

Chest X-ray is widely used to evaluate the proper placement of catheters, as incorrect placement can lead to serious complications. Misplacement of an endotracheal tube (ETT) can result in conditions such as hypoxemia, pneumothorax, and pneumonia (Lakhani et al., 2021), while a misplaced nasogastric tube (NGT) can cause respiratory failure and even death (Singh et al., 2019). Studies have shown that the estimated misplacement rate of ETT and NGT catheters can be as high as 28% and 15%, respectively (Yi et al., 2020). With the advancement of computer vision and deep learning, computer-aided detection (CAD) methods have been developed to assist radiologists in quickly and accurately detecting catheter misplacement (Aryal and Yahyasoltani, 2021; Elaanba et al., 2021; Lakhani et al., 2021; Lakhani, 2017; Singh et al., 2019). However, for CAD to be clinically useful, it is important to also consider model explainability. This has led to the development of semantic segmentation methods for catheter placement on chest X-ray images, which can provide detailed information about the location and shape of the catheter (Sullivan et al., 2020; Gherardini et al., 2020).

However, obtaining large datasets with pixel-wise labels is often difficult as it requires a lot of resources and expert knowledge to perform the annotation of the images. Additionally,

© 2023 Anonymous.

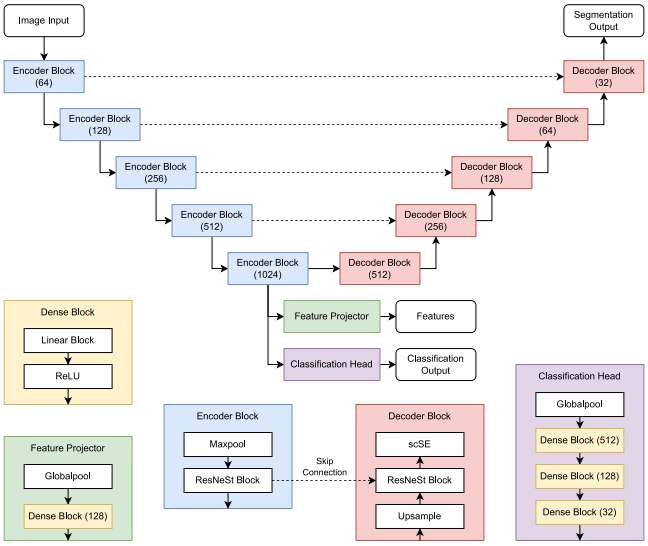

Figure 1: U-Net with feature projector and classification head.

class imbalance is a common problem in catheter placement classification, as misplaced catheters are less common than correctly placed catheters. This makes it challenging for a model to accurately segment misplaced catheters and predict the absence or misplacement of catheters. Semi-supervised contrastive learning techniques have been used to address these issues by utilizing large unannotated datasets (Hu et al., 2021).

At the same time, realizing that the catheter is a physically connected object, we can use its expected topology as a prior to compute a topology loss. This loss does not require labeling as the catheter's topology is assumed to be the same across all images. Therefore, we propose a semi-supervised learning method that combines a contrastive loss with a topology loss to improve the performance of catheter placement classification and segmentation.

We developed an algorithm that simultaneously performs catheter segmentation and misplacement prediction using a U-Net architecture with an additional classification head. This model can predict the misplacement of the ETT or NGT catheters, in addition to segmenting the catheters and relevant anatomical landmarks. This allows our model to provide not only detailed information about the catheter location but also its malposition state, and hence make the model more useful in clinical practice and explainable.

Our contributions in this work can be summarized as follows:

- Developed a catheter segmentation and misplacement prediction deep learning model;
- Validated the proposed algorithm on external testing set;
- Combined contrastive and topology losses for semi-supervised learning;
- Modified topology loss to constraint on a range of topological features.

2. Methods

2.1. Model

The model we use is based on U-Net (Ronneberger et al., 2015), a popular fully convolutional neural network architecture for biomedical image segmentation. U-Net contains an encoder that extracts features of different spatial resolutions, and uses skip connections to pass feature maps to a decoder that generates a segmentation mask through pixel-wise classification. The encoder in our model uses ResNeSt-50 (Zhang et al., 2020) as its backbone, which implements a channel-wise attention mechanism to capture cross-feature interactions and learn diverse representations. Additionally, we apply Spatial and Channel Squeeze and Excitation (scSE) (Roy et al., 2019) to the decoder blocks, which aggregates spatial information and feature information to enhance the segmentation performance. To perform both segmentation and misplacement classification, we add a feature projector and a classification head after the encoder, in addition to the U-Net architecture for segmentation. The feature projector is used to convert the features extracted by the encoder into a feature representation that can be used in the contrastive learning. The overall architecture of the model is illustrated in Fig. 1.

2.2. Contrastive Loss

Contrastive loss is a loss function used in self-supervised contrastive learning, which aims to learn visual representations from unlabeled images. The framework we use is the SimCLR framework proposed by (Chen et al., 2020). The approach is to train the model to differentiate between different views of the same data. This allows the model to learn representations of the data that can be used for downstream tasks such as segmentation and classification.

The implementation is as follows, for each input batch $B = \{x_1, x_2, ..., x_b\}$, where x_i represents each input image, every image is performed augmentations twice to obtain $B' = \{x_1, x_2, ..., x_{2b}\}$. Let $\sigma(i)$ be the index of image that is derived from the same image as index i. Then denote the encoder as $E(\cdot)$ and the feature projector as $g(\cdot)$. The contrastive loss of a batch is defined as

$$\mathrm{Contr}(B') = -\frac{1}{2b}\sum_{i=1}^{2b}\log\frac{\exp(\mathrm{sim}(z_i, z_{\sigma(i)})/\tau)}{\sum_{k=1}^{2N}\mathbb{1}_{[k\neq i]}\exp(\mathrm{sim}(z_i, z_k)/\tau)} \tag{1}$$

where τ is the temperature, $z_i = g(E(x_i))$ and $\mathrm{sim}(u, v) = \frac{u^T v}{||u||\cdot||v||}$ which is the cosine similarity. The augmentations we are using is the composition of random contrast adjustment, Gaussian noise, dropout, random grid distortion and random affine transformation.

2.3. Topology Loss

The topology loss (Clough et al., 2020) is a segmentation loss that is based on persistent homology, a mathematical framework for the study of topological features of shapes. The topology loss does not require segmentation labels but instead uses topology priors, such as Betti numbers, to impose constraints on the segmentation output. The Betti numbers are topological invariants that describe topological properties such as the number of connected

components and the number of holes in a shape. The loss function consists of two components, the level set and the barcode diagram. The level set is the collection of segmentation results at different probability thresholds, and the barcode diagram encodes information about the topological changes that occur as the threshold is varied.

Notice that when the threshold is below zero, the entire segmentation would be considered one connected component. In theory, this would lead to an infinitely long bar in the barcode diagram. Hence, the authors of the topology loss suggested that this bar can be cut off at a threshold of 0 (Clough et al., 2020). We implemented the topology loss such that the bar is cut off at zero. By doing this, we can impose the segmentation to be either zero or one connected component.

We also modified the loss such that the target can be applied on a range of acceptable Betti numbers, instead of an exact value. For example, in our case the catheter might or might not be present in the image, thus the number of connected components can be zero or one. To be able to use this loss, we can modify its implementation such that we only penalize when the output is beyond our target range. Thus, the topology loss of an segmentation output is defined as

$$\text{Topo}(\hat{y}) = \sum_{k} \left(\sum_{l=1}^{\beta_{k,a}} (1 - |B_{l,k}(\hat{y})|)^2 + \sum_{l=\beta_{k,b}+1}^{\infty} |B_{l,k}(\hat{y})|^2 \right) \tag{2}$$

where $|B_{l,k}(\hat{y})|$ is the length of l-th longest barcode of dimension k of output $\hat{y}$. $\beta_{k,a}$ and $\beta_{k,b}$ are respectively the lower and upper bound of the acceptable Betti number of dimension k. Notice that the second term has infinitely many terms, but in practice there can only be a finite number of topology features in output of a finite image. In fact, most barcodes will be zero, hence we set the the maximum number of terms to be 20.

In our task, we apply the topology loss to the segmentation prediction of catheters, where it can be at most one connected component. We also did not impose a constraint on the number of holes present in the segmentation, since it is possible for catheter to curl and form loops. It is important to note that computing this topological loss is computationally expensive and is currently only implemented on CPU. To mitigate this issue, we downscale the output by max-pooling to a resolution of 16×16, which reduces the number of pixels, resulting in a less expensive computation. However, this also ensures that the overall structure of the segmentation prediction is still constrained.

2.4. Semi-supervised Learning

Figure 2 shows the semi-supervised training pipeline that we proposed. Noticing that both the contrastive loss and the topology loss do not require labels to be computed, we combine both method for semi-supervised learning. The total loss for a batch is defined as

$$\text{Total Loss} = \omega_D \text{Dice}_l + \omega_F \text{Focal}_l + \omega_T \text{Topo}_l + \alpha_u \omega_T \text{Topo}_u + \omega_C \text{Contr}_{l,u} \tag{3}$$

In the above equation, Dice, Focal, Topo, and Contr represent the dice loss, focal loss, topology loss, and contrastive loss, respectively. The subscript u denotes that the loss is applied to unlabeled data, while the subscript l denotes that the loss is applied to labeled data. Note that contrastive loss is applied on the combined batch of labeled and unlabeled

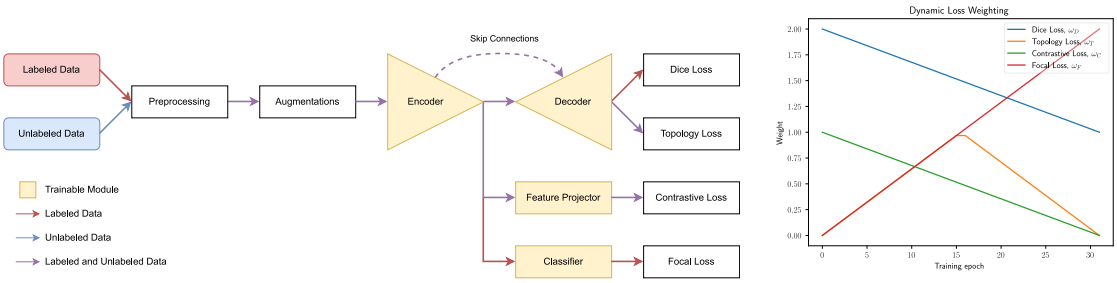

Figure 2: The semi-supervised training pipeline (left) and dynamic loss weighting (right).

data. ω_D, ω_F, ω_T, ω_C are the weights for each loss term. α_u is a factor used to account for the ratio of unlabeled and labeled data in the batch during training.

To effectively train a diverse task that requires aggregating multiple loss functions, a dynamic weighting mechanism is implemented during training. In the initial stages of training, the contrastive and segmentation losses are given higher weight as the visual representation and segmentation need to be optimized first as they are crucial for the classification and topology tasks. As the training progresses, the focus shifts to optimizing the topological properties of the segmentation output. Once the segmentation is well-optimized, the model can be trained on the classification task more effectively.

3. Experiments and Results

3.1. Dataset

3.1.1. *****

We collected 7,378 supine portable chest X-rays from the ***** (*****) which are labeled with the misplacement state of NGT and ETT, as well as segmentation annotation for the catheters and relevant anatomical landmarks. This dataset includes a training set and a testing set. The training set (n=5,767) was collected from the ***** PACS database through keyword search of the radiology information system, and it includes data collected from the emergency department or ICUs between 2015 and 2019. The testing set (n=1,611) was randomly sampled from the emergency department and is a mixture of data from 2020 from ***** and data from 2015-2020 from the ***** Branch to test the chronological and geographical generalizability of the model. Each chest X-ray is labeled as Abnormal (incorrect placement), Normal (correct placement), or NA (catheter absence) for each catheter. Figure 3 shows the distribution of classification label.

Figure 4 is a sample of chest X-ray and its segmentation annotations. The segmentation annotation includes labels for ETT (light brown), NGT (light green), Lung (yellow), Trachea (red), Diaphragm (orange), as well as three additional channels for ETT tip (dark brown), NGT tip (dark green) and Carina (pink), that was extracted from the segmentation mask of ETT, NGT, and Trachea.

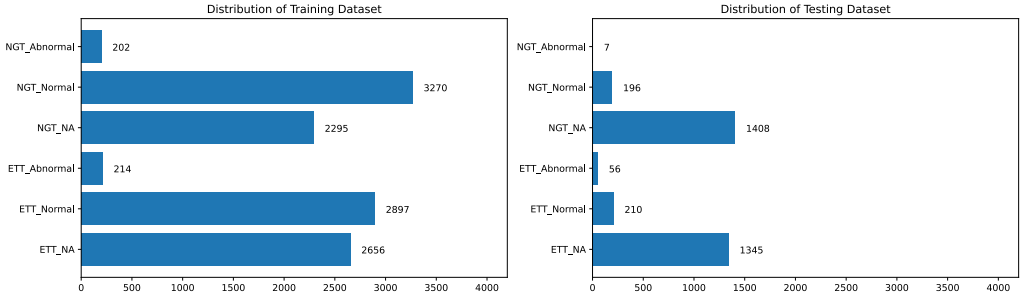

Figure 3: The distribution of catheter misplacement of ***** dataset.

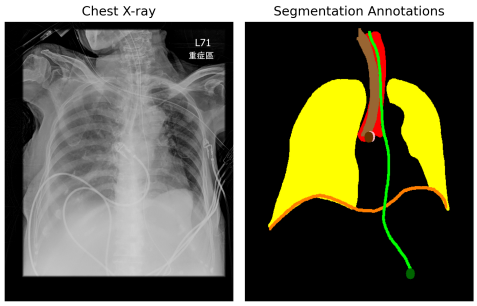

Figure 4: Sample of chest X-ray and its segmentation annotations.

3.1.2. CLiP

This Catheter and Line Position (CLiP) dataset is provided by (Tang et al., 2021). It contains 30,083 chest X-rays that includes data with misplaced catheter. It contains misplacement label of NGT and ETT, as well as their segmentation mask. In our work, this dataset is used as unlabeled data for semi-supervised learning.

3.2. Training

In the experiment, the data is split into training and validation sets with a ratio of 4 : 1 using a stratified sampling method, ensuring the same proportion of each class. Images with catheter malposition were oversampled during training to balance the number across each class.

As preprocessing, each input image is resized to 512 × 512 before applying contrast-limited adaptive histogram equalization (CLAHE) (Pizer et al., 1990) to enhance the image. Nine different parameters pairs were used for the CLAHE, by taking 0.02, 0.05, 0.1 of the image size as the kernel size and 0.02, 0.05, 0.1 as the clip limit. The image is then processed with CLAHE nine times and the results are stacked into a nine-channel array to improve the visibility.

During the training process, the batch size was 36, with the proportion of labeled vs unlabeled data of 3:1, the AdamW optimizer (Loshchilov and Hutter, 2017) was used, and the learning rate was 3e-4, scheduled using Cosine Annealing with Warm Restarts (Loshchilov

Table 1: Performance of different methods.

Previous Studies	Dice (95% CI)	AUC (95% CI)
Fully-supervised (Sullivan et al., 2020)	0.74	-
Fully-supervised (Elaanba et al., 2021)	-	0.80
Fully-supervised (Lakhani, 2017)	-	0.81
Fully-supervised (Singh et al., 2019)	-	0.87 (0.80, 0.94)
Synthetic data (Gherardini et al., 2020)	0.55	-
Synthetic data (Aryal and Yahyasoltani, 2021)	-	0.96
Ours		
Fully-supervised (baseline)	0.520 (0.515, 0.526)	0.979 (0.968, 0.988)
Semi-supervised TL	0.372 (0.366, 0.377)	0.974 (0.965, 0.983)
Semi-supervised CL	0.562* (0.556, 0.568)	0.978 (0.955, 0.979)
Semi-supervised CL + TL	0.598* (0.592, 0.604)	0.977 (0.968, 0.986)

TL = Topology loss, CL = Contrastive loss
* Statistically significant over baseline with p-value < 0.05

and Hutter, 2016). To have more stable and realistic results of the segmentation task, the output and target segmentation masks were applied with dilation using a 3×3 square kernel. Otherwise, the line-like shape of the segmentation could cause the loss and metrics to be oversensitive.

The segmentation loss used in the experiment is the spatial-weighted dice loss, which is a modification of the dice loss (Milletari et al., 2016) that assigns more weight to pixels closer to the point of interests (POIs), namely NGT tip for NGT, ETT tip for ETT, and Carina for Trachea. The classification loss used is focal loss (Lin et al., 2017), with $\gamma = 2$ and $\alpha = 3$ for misplaced catheter. This loss also includes label smoothing (Müller et al., 2019) of 0.1.

3.3. Results

The performance of our proposed method is evaluated on the ***** testing set, using metrics including Dice coefficient and area under the receiver operating characteristic curve (AUC). We also compute 95% confidence intervals and p-values over baseline using bootstrapping (with n=1000). The Dice coefficient is used to measure the agreement between the predicted and target segmentation masks. The predicted mask is thresholded at 0.5, and the result is an average over all predicted classes. AUC is used to evaluate the classification performance. It is computed with one-vs-rest for NA, Normal, and Abnormal and then take the macro average. The result is the mean AUC over all catheters. The results are shown in Table 1. For result visualizations, see Appendix A.

4. Discussion

Previous studies on catheter segmentation have used supervised learning (Sullivan et al., 2020), and have achieved Dice scores of 0.74, and for catheter misplacement prediction, AUC values above 0.8 (Elaanba et al., 2021; Lakhani, 2017; Singh et al., 2019). However, these studies have been limited by the availability of labeled data. To address this problem, some studies have used synthetic data methods, such as imposing predefined segmentation

labels on X-ray images (Gherardini et al., 2020), or using generative adversarial networks (GANs) to synthesize catheters in X-ray images (Aryal and Yahyasoltani, 2021), achieving Dice scores of 0.55 and AUC values of 0.96. It is important to note that these mentioned studies, although similar in nature, have slightly different tasks and datasets. Therefore, direct comparison between them is not possible and they are only listed for reference.

In contrast, our proposed method utilizes a semi-supervised learning approach by using a contrastive loss to train image representations and a topology loss to train segmentation with topological priors, which allows the model to learn from both labeled and unlabeled data. This approach improves the baseline performance of the segmentation significantly and reaches a Dice score of 0.598, while maintaining the performance for misplacement prediction at AUC 0.977. The effect of using only topology loss and contrastive loss on classification performance is limited, as no additional information about the classification labels is provided during the semi-supervised learning process.

From the results shown in Table 1, it is clear that using the topology loss alone decreases the performance of Dice score from 0.520 to 0.372. However, when combined with contrastive loss, the performance improves to 0.589, which is higher than using contrastive loss alone. This can be attributed to the fact that topology loss is only effective when the initial segmentation results are already reasonable. On the other hand, using contrastive loss allows the model to learn visual representations from unlabeled data, thus improving the performance of segmentation from 0.520 to 0.562. Therefore, by adding topology loss on top of contrastive loss, the model can improve the initial segmentation performance from contrastive loss and then apply topology loss to restrict the number of connected components more effectively, resulting in an improved performance of 0.598. This result is statistically significant over the baseline with a p-value of less than 0.05.

5. Conclusion

In summary, this study proposed a semi-supervised learning method that combines contrastive loss and topology loss to improve the performance of catheter segmentation and misplacement prediction on chest X-ray images. The method utilizes a U-Net architecture with a classification head and learns from both labeled and unlabeled data by using a contrastive loss to train image representations and a topology loss to train segmentation with topological priors. The results of the study show that the proposed method improves the baseline performance of segmentation significantly and reaches a average Dice score of 0.598, while maintaining the performance for misplacement prediction at average AUC 0.977. Our proposed method of combining contrastive loss and topology loss for semi-supervised learning can be extended to other tasks to potentially improve the performance of segmentation, as long as additional unlabeled data is available and a topological prior is known.

References

Milan Aryal and Nasim Yahyasoltani. Identifying catheter and line position in chest x-rays using gans. In *2021 20th IEEE International Conference on Machine Learning and Applications (ICMLA)*, pages 122–127, 2021. doi: 10.1109/ICMLA52953.2021.00027.

Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey Hinton. A simple framework for contrastive learning of visual representations. In Hal Daumé III and Aarti Singh, editors, *Proceedings of the 37th International Conference on Machine Learning*, volume 119 of *Proceedings of Machine Learning Research*, pages 1597–1607. PMLR, 2020. URL https://proceedings.mlr.press/v119/chen20j.html.

James Clough, Nicholas Byrne, Ilkay Oksuz, Veronika Zimmer, Julia Schnabel, and Andrew King. A topological loss function for deep-learning based image segmentation using persistent homology. *IEEE transactions on pattern analysis and machine intelligence*, PP, 2020. doi: 10.1109/TPAMI.2020.3013679.

Abdelfettah Elaanba, Mohammed Ridouani, and Larbi Hassouni. Automatic detection using deep convolutional neural networks for 11 abnormal positioning of tubes and catheters in chest x-ray images. In *2021 IEEE World AI IoT Congress (AIIoT)*, pages 0007–0012, 2021. doi: 10.1109/AIIoT52608.2021.9454205.

Marta Gherardini, Evangelos Mazomenos, Arianna Menciassi, and Danail Stoyanov. Catheter segmentation in x-ray fluoroscopy using synthetic data and transfer learning with light u-nets. *Computer Methods and Programs in Biomedicine*, 192:105420, 2020. ISSN 0169-2607. doi: https://doi.org/10.1016/j.cmpb.2020.105420.

Xinrong Hu, Dewen Zeng, Xiaowei Xu, and Yiyu Shi. Semi-supervised contrastive learning for label-efficient medical image segmentation. In Marleen de Bruijne, Philippe C. Cattin, Stéphane Cotin, Nicolas Padoy, Stefanie Speidel, Yefeng Zheng, and Caroline Essert, editors, *Medical Image Computing and Computer Assisted Intervention – MICCAI 2021*, pages 481–490, Cham, 2021. Springer International Publishing. ISBN 978-3-030-87196-3.

P. Lakhani. Deep convolutional neural networks for endotracheal tube position and x-ray image classification: Challenges and opportunities. *J Digit Imaging*, 30(4):460–468, 2017. doi: 10.1007/s10278-017-9980-7. PMID: 28600640.

Paras Lakhani, Adam Flanders, and Richard Gorniak. Endotracheal tube position assessment on chest radiographs using deep learning. *Radiology: Artificial Intelligence*, 3(1): e200026, 2021. doi: 10.1148/ryai.2020200026. PMID: 33937852.

Tsung-Yi Lin, Priya Goyal, Ross Girshick, Kaiming He, and Piotr Dollár. Focal loss for dense object detection. In *2017 IEEE International Conference on Computer Vision (ICCV)*, pages 2999–3007, 2017. doi: 10.1109/ICCV.2017.324.

Ilya Loshchilov and Frank Hutter. SGDR: stochastic gradient descent with restarts. *CoRR*, abs/1608.03983, 2016. URL http://arxiv.org/abs/1608.03983.

Ilya Loshchilov and Frank Hutter. Fixing weight decay regularization in adam. *CoRR*, abs/1711.05101, 2017. URL http://arxiv.org/abs/1711.05101.

Fausto Milletari, Nassir Navab, and Seyed-Ahmad Ahmadi. V-net: Fully convolutional neural networks for volumetric medical image segmentation. *CoRR*, abs/1606.04797, 2016. URL http://arxiv.org/abs/1606.04797.

Rafael Müller, Simon Kornblith, and Geoffrey E Hinton. When does label smoothing help? In H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc., 2019. URL https://proceedings.neurips.cc/paper/2019/file/f1748d6b0fd9d439f71450117eba2725-Paper.pdf.

S.M. Pizer, R.E. Johnston, J.P. Ericksen, B.C. Yankaskas, and K.E. Muller. Contrast-limited adaptive histogram equalization: speed and effectiveness. In *[1990] Proceedings of the First Conference on Visualization in Biomedical Computing*, pages 337–345, 1990. doi: 10.1109/VBC.1990.109340.

Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation. In *International Conference on Medical image computing and computer-assisted intervention*, pages 234–241, 2015.

Abhijit Guha Roy, Nassir Navab, and Christian Wachinger. Recalibrating fully convolutional networks with spatial and channel "squeeze and excitation" blocks. *IEEE Transactions on Medical Imaging*, 38(2):540–549, 2019. doi: 10.1109/TMI.2018.2867261.

V. Singh, V. Danda, R. Gorniak, A. Flanders, and P. Lakhani. Assessment of critical feeding tube malpositions on radiographs using deep learning. *J Digit Imaging*, 32(4):651–655, 2019. doi: 10.1007/s10278-019-00229-9. PMID: 31073816.

Ryan Sullivan, Greg Holste, Jonathan Burkow, and Adam Alessio. Deep learning methods for segmentation of lines in pediatric chest radiographs. In Horst K. Hahn and Maciej A. Mazurowski, editors, *Medical Imaging 2020: Computer-Aided Diagnosis*, volume 11314, pages 577 – 583. International Society for Optics and Photonics, SPIE, 2020. doi: 10.1117/12.2550686.

Jennifer S. N. Tang, Jarrel C. Y. Seah, Adil Zia, Jay Gajera, Richard N. Schlegel, Aaron J. N. Wong, Dayu Gai, Shu Su, Tony Bose, Marcus L. Kok, Alex Jarema, George N. Harisis, Chris-Tin Cheng, Helen Kavnoudias, Wayland Wang, Anouk Stein, George Shih, Frank Gaillard, Andrew Dixon, and Meng Law. Clip, catheter and line position dataset. *Scientific Data*, 8(1):285, 2021. ISSN 2052-4463. doi: 10.1038/s41597-021-01066-8.

Xin Yi, Scott J. Adams, Robert D. E. Henderson, and Paul Babyn. Computer-aided assessment of catheters and tubes on radiographs: How good is artificial intelligence for assessment? *Radiology: Artificial Intelligence*, 2(1):e190082, 2020. doi: 10.1148/ryai.2020190082. PMID: 33937813.

Hang Zhang, Chongruo Wu, Zhongyue Zhang, Yi Zhu, Zhi Zhang, Haibin Lin, Yue Sun, Tong He, Jonas Muller, R. Manmatha, Mu Li, and Alexander Smola. Resnest: Split-attention networks. *arXiv preprint arXiv:2004.08955*, 2020.

Appendix A. Result Visualizations

Result visualizations of different experiments are presented in Fig. 5.

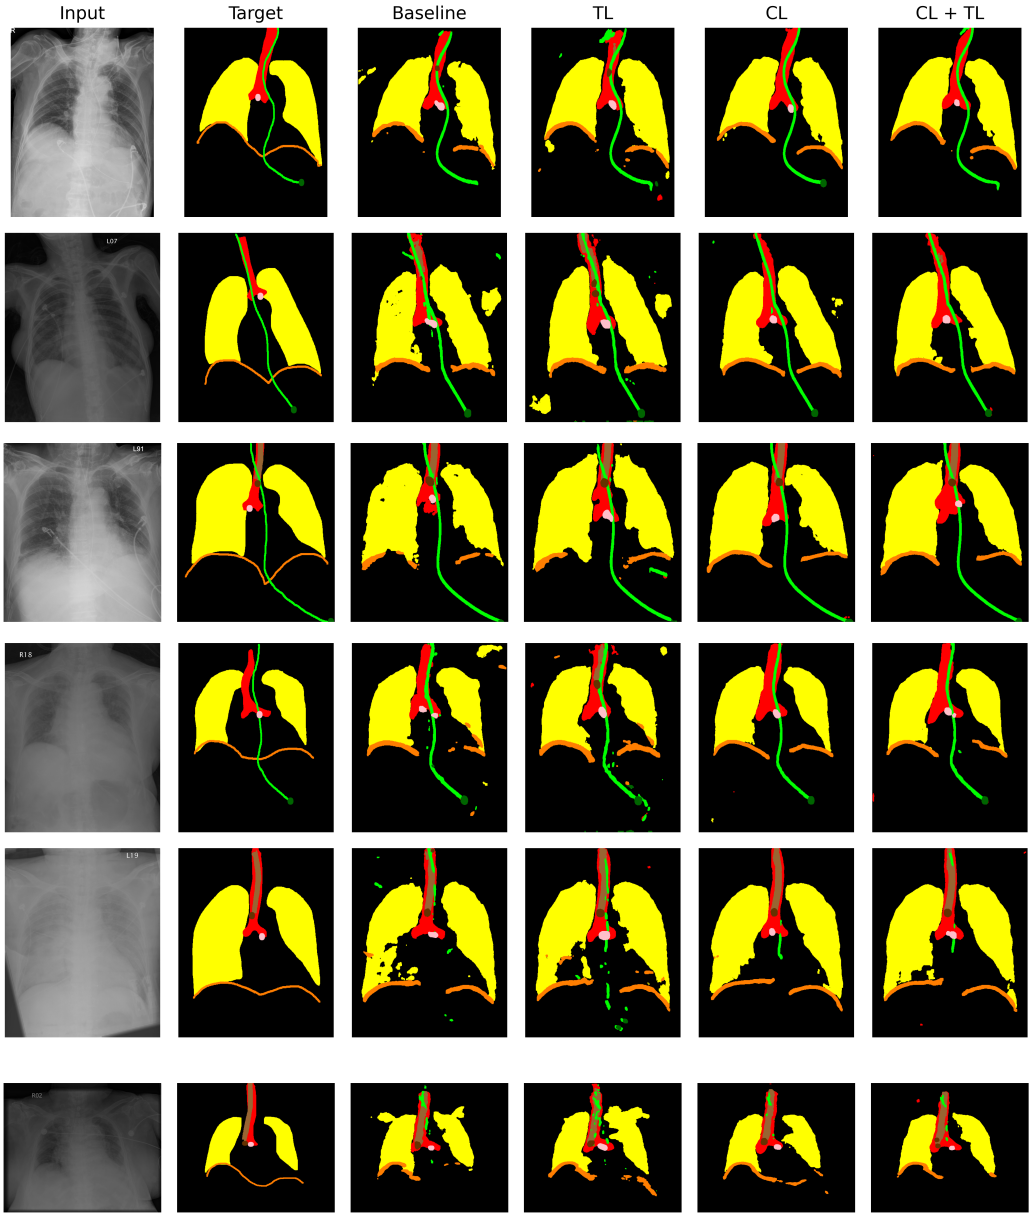

Figure 5: Result visualizations of different experiments. From left to right are input chest X-rays, segmentation targets, and sample predictions of Fully-supervised, Semi-supervised TL, Semi-supervised CL, and Semi-supervised CL + TL experiments. From top to bottom are six different samples.