

Detection of Peri-Pancreatic Edema using Deep Learning and Radiomics Techniques

Ziliang Hong³, Debesh Jha¹, Koushik Biswas¹, Zheyuan Zhang¹, Yury Velichko¹, Cemal Yazici³, Temel Tirkes⁴, Amir Borhani¹, Baris Turkbey², Alpaly Medetalibeyoglu¹, Gorkem Durak¹, Ulas Bagci¹

Abstract—Pancreatitis is a major public health issue worldwide; studies show an increase in the number of people experiencing pancreatitis. Identifying peri-pancreatic edema is a pivotal indicator for identifying disease progression and prognosis, emphasizing the critical need for accurate detection and assessment in pancreatitis diagnosis and management. This study introduces a novel CT dataset sourced from 255 patients with pancreatic diseases, featuring annotated pancreas segmentation masks and corresponding diagnostic labels for peri-pancreatic edema condition. With the novel dataset, we first evaluate the efficacy of the *LinTransUNet* model, a linear Transformer based segmentation algorithm, to segment the pancreas accurately from CT imaging data. Then, we use segmented pancreas regions with two distinctive machine learning classifiers to identify existence of peri-pancreatic edema: deep learning-based models and a radiomics-based eXtreme Gradient Boosting (XGBoost). The *LinTransUNet* achieved promising results, with a dice coefficient of 80.85%, and mIoU of 68.73%. Among the nine benchmarked classification models for peri-pancreatic edema detection, *Swin-Tiny* transformer model demonstrated the highest recall of 98.85 ± 0.42 and precision of 98.38 ± 0.17 . Comparatively, the radiomics-based XGBoost model achieved an accuracy of 79.61 ± 4.04 and recall of 91.05 ± 3.28 , showcasing its potential as a supplementary diagnostic tool given its rapid processing speed and reduced training time. To our knowledge, this is the first study aiming to detect peri-pancreatic edema automatically. We propose to use modern deep learning architectures and radiomics together and created a benchmarking for the first time for this particular problem, impacting clinical evaluation of pancreatitis, specifically detecting peri-pancreatic edema. Our code is available <https://github.com/NUBagciLab/Peri-Pancreatic-Edema-Detection>.

I. INTRODUCTION

Pancreatitis poses a significant global healthcare concern due to its potential mortality risk [17], [19]. The presence of peri-pancreatic edema serves as a crucial indicator of disease progression and prognosis. Thus, accurate peri-pancreatic edema detection and assessment are pivotal in early pancreatitis diagnosis and management. This is important because it can help determine the severity and progression of the condition, vital for improving patient prognosis. Despite advancements in the field of medical imaging with several innovative deep-learning approaches [2], [15], [9], [1], the precise quantification and analysis of peri-pancreatic edema

is still an unsolved problem. Thus, there is an emergent need for developing techniques to solve the problem of peri-pancreatic edema detection.

Deep learning is now a de-facto standard to solve the challenges related to many medical imaging applications. However, applying these advanced deep learning (DL) techniques in the specific context of pancreatitis and peri-pancreatic edema detection has yet to be explored. As DL-based techniques could facilitate early diagnosis and treatment planning, this study aims to investigate the effectiveness of a DL-based segmentation algorithm, specifically *LinTransUNet*, in accurately segmenting the pancreas from CT imaging data [18]. Subsequently, the extracted segmentation regions (ROI) are designed as training data for a Vision Transformer (ViT) model and radiomics-based XGBoost model in our experiments for peri-pancreatic edema prediction [5], [13]. This research serves as a complementary tool to clinical diagnosis of pancreatitis by exploring the application of deep learning and radiomics models in pancreatic CT image analysis. Furthermore, the findings of this study hold valuable insights into the utilization of deep learning techniques for efficient medical data labeling. In summary, the contributions of this paper are as follows:

(1) We are not aware of any machine learning or deep learning methodology automatically detecting peri-pancreatic edema in CT scans. Therefore, our method serves as a strong baseline for peri-pancreatic edema detection even with conventional machine learning approaches. This is the first-ever application of pancreatic edema detection, to our best of knowledge.

(2) We provide a unique dataset for the public, comprising CT imaging data from 255 patients with detailed pancreas segmentation masks. This comprehensive dataset will serve as a valuable resource for the exploration and advancement of deep learning techniques in pancreatic research in general, offering a rich source of information for future research and clinical applications. The public links will be available upon acceptance of the paper.

(3) We introduce two distinct classifiers: a deep learning-based and a radiomics-based classifier. The utility of the generated pancreas segmentation mask supports the radiomics analysis. These classifiers, built upon the extracted ROIs from our segmentation, serve as efficient tools for assisting medical professionals in diagnosing pancreatic diseases. Radiomics is a strong tool when the data is small. Herein, we combine the complementary power of both deep learning

*This work is supported by the NIH funding: R01-CA246704, R01-CA240639, U01-DK127384-02S1, and U01-CA268808.

¹D. Jha, K. Biswas, A. Bourhani, G. Durak, Z. Zhang, and U. Bagci are with Machine & Hybrid Intelligence Lab, Northwestern University. ²B. Turkbey is with National Cancer Institute. ³Z. Hong and C. Yazici are with University of Illinois at Chicago. ⁴T. Tirkes is with Indiana University.

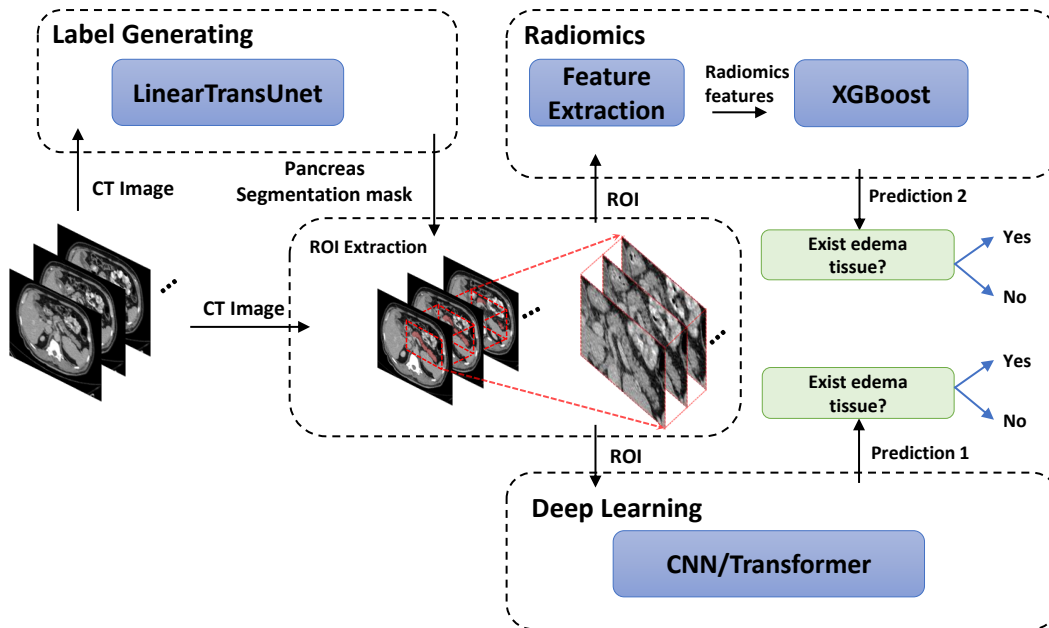


Fig. 1. Proposed workflow for detecting peri-pancreatic edema from CT images using two machine learning approaches: *radiomics* and *deep learning*. The process begins with inputting CT images into the *LinearTransUnet* [18] segmentor for pancreas label generation. Segmented pancreases are in two complementary paths. In the first path, radiomics techniques are applied to the ROI (segmented pancreases) to extract features that are then used by an eXtreme Gradient Boosting (XGBoost) model to make a prediction (Prediction 2) regarding the presence of edema tissue. The second path utilizes the same ROIs with a deep learning-based model (CNN & Transformers) for diagnostic prediction.

and radiomics to improve the detection of pancreatic peri-pancreatic edema. Deep learning and radiomics models as well as benchmarking methods will be available to public along with the data upon publication of the manuscript.

II. METHODOLOGY

Deep Learning Classification. Once pancreases are segmented, regions tightly enclosing the pancreases can be used for automatic prediction of the peri-edema status. While we are using *LinearTransUnet* [18] for the segmentation, one may use other segmentation strategies including semi-automated or manual labelings too. *LinearTransUnet* is chosen as our segmentation method because of its outstanding performance in image segmentation.

We use a rich set of architecture family for benchmarking classification results. For convolutional neural networks (CNN) based comparison, residual network (ResNet) [6], densely connected convolutional networks (DenseNet) [8], MobileNetV3 [7], MobileNetV2 [14], and EfficientNet [4] were included in the experiment. Representative transformer models like vision transformer (ViT) [5] and Swin transformer (Swin) [12] [11] were also analyzed for comprehensive study. Models are built strictly following their corresponding papers. The number of input channels is equal to 1 since input images are in grayscale. The trainable data was standardized to a consistent format and resized uniformly to 224×224 using linear interpolation. For deep learning models, we trained the model for 100 epochs in the batch size of 128 with 5-fold cross-validation, ensuring the robustness of the evaluations. Leveraging Stochastic Gradient

Descent (SGD) with a momentum coefficient of 0.9 and original learning rate of $1e^{-5}$ as an optimizer, and cross-entropy as the loss function, respectively, the model aimed to optimize parameters and mitigate overfitting. The entire training process was executed using the PyTorch framework on an NVIDIA RTX A6000 GPU.

Radiomics based classification. In radiomics analysis, the necessity of the large quantities extraction of quantitative image features has been demonstrated in reports by Lambin [10] and Yao [17]. This study covers 107 distinct features extracted from CT images based on the generated masks. These features contain the following kinds: First Order Features, Shape Features, Gray Level Co-occurrence Matrix (GLCM) Features, Gray Level Size Zone Matrix (GLSZM) Features, Gray Level Run Length Matrix (GLRLM) Features, Neighbouring Gray Tone Difference Matrix (NGTDM) Features, Gray Level Dependence Matrix (GLDM) Features. XGboost was considered the key approach to build the radiomics-based classifier according to its outstanding performance [3]. Optimal hyper-parameters were identified by adopting the grid search. This research set the following hyper-parameters: estimators = 3, maximum depth = 2.

III. EXPERIMENTS AND RESULTS

A. Dataset Overview

We collected CT scans of 255 pancreatitis patients with IRB approval from Istanbul University Capa School of Medicine. There were two classes in the dataset: **179** pancreatitis patients **with peri-pancreatic edema** and **76** pancreatitis patients **without peri-pancreatic edema**. Each




Fig. 2. The upper row showcases the original CT scans, while the bottom row displays the extracted pancreas ROI. The images are presented in three different planes: axial, coronal, and sagittal, arranged from left to right.

of the images had a corresponding standard pancreas mask and a label for disease class. Images with peri-pancreatic edema tissue were categorized as positive samples, denoted by label 1, other images are denoted with label 0. All standard pancreas segmentation masks were annotated by a team of expert radiologists (N=2) on consensus. After annotation, a careful evaluation was made to ensure the accuracy and reliability of the annotations. The standard pancreas segmentation masks were then used to evaluate the pancreas mask-generating model. Our dataset is available at <https://osf.io/wyth7/>.

B. Preprocessing

The noise of medical instruments or motion blur can degrade the quality of medical images. Inspired by Yao’s method [16], our first step was to perform reorientation. It is important, especially for cases where the reconstructed scan is not in the same orientation as with a mask. This is particularly important for the reproducibility of radiomics studies. Secondly, we applied conventional denoising algorithms (Gaussian denoising filter) to smooth out the noise. Once the denoising was performed, we used LinTransUNet [18] to segment the pancreases.

We employed a small extension for the ROI enclosing the pancreas region. We used 10% scaling in all axis to ensure the completeness of the pancreas region of over image. An extended ROI example is shown in Fig. 2.

C. Segmentation Results

For assessing the segmentation performance of LinTransUNet [18], the Dice score, mIoU, precision, and recall were employed. With LinTransUNet, we obtained a Dice score of 80.85%, mIoU of 68.73%, precision of 76.21%, and recall of 87.65% as shown in Table I.

D. Classification

The evaluation of the classifiers on our pancreas dataset relied on accuracy, precision, and recall as key metrics to provide further insights into the model’s classifying of pancreatic and peri-pancreatic edema tissue. The outcomes obtained from 5-fold cross-validation are tabulated in Table I. As described in Table I, Swin-tiny owns the highest accuracy,

TABLE I
PERFORMANCE COMPARISON OF CNN AND TRANSFORMER BASED
APPROACH FOR PERI-PANCREATIC EDEMA DETECTION.

Backbone	Accuracy (%)	Precision (%)	Recall (%)
EfficientNet B0 [4]	79.16 (± 0.38)	98.53 (± 0.52)	78.60 (± 0.55)
MobileNetV2 [14]	81.31 (± 0.58)	97.56 (± 0.54)	80.97 (± 0.72)
MobileNetV3-s [7]	85.92 (± 1.38)	96.91 (± 0.37)	85.73 (± 1.40)
ViT-tiny [5]	93.20 (± 1.62)	97.11 (± 0.30)	93.85 (± 2.23)
ResNet-50 [6]	95.09 (± 0.39)	97.03 (± 1.39)	96.22 (± 0.63)
DenseNet-121 [8]	96.41 (± 0.34)	98.21 (± 0.61)	96.96 (± 0.32)
SwinV2-tiny [11]	97.27 (± 0.52)	98.60 (± 0.80)	97.71 (± 0.45)
ResNet-34 [6]	97.63 (± 0.20)	98.54 (± 0.52)	98.25 (± 0.43)
Swin-tiny [12]	97.95 (± 0.40)	98.85 (± 0.42)	98.38 (± 0.17)

TABLE II
RESULT OF RADIOMICS-BASED XGBOOST [3] MODEL ON TEST DATA.

Fold	Accuracy (%)	Precision (%)	Recall (%)
1	78.43	80.49	91.67
2	88.24	87.50	97.22
3	76.47	80.00	88.89
4	78.43	82.05	88.89
5	76.47	79.49	88.57
Average	79.61 (± 4.04)	81.91 (± 2.93)	91.05 (± 3.28)

precision, and recall. ResNet-34 also indicated a very high value in all metrics, and these values are close to those of Swin-tiny. SwinV2-tiny and DenseNet-121 also show relatively good results. The accuracy of the two models reached 96%. ResNet-50 and ViT-tiny underperform in the results, with 95% and 93% accuracy, respectively. The above models performed well in this experiment, with accuracy exceeding 90%, precision exceeding 97%, and recall exceeding 90%.

Table I demonstrates that MobilenetV2 MobileNetV3-s and EfficientNet B0 were relatively poor compared to others but still considerable in performance. MobilenetV2’s accuracy was less than 86%. Besides, its recall decreased below 86%. A similar performance decrease was also shown in MobileNetV3-s and EfficientNet B0 and the recall drops to 78% in the results of the EfficientNet B0, which indicates an erroneous diagnosis rate of nearly 22% for patients with peri-pancreatic edema tissue. In contrast, the precision results were notably robust, each deep learning model exhibits a precision exceeding 96%.

With the same evaluation metrics as the deep learning model, the results of the radiomics-based XGBoost model are presented in Table II. From this data, we can observe that the XGBoost model underperformed with both metrics compared with the deep learning model in general. XGBoost model has 79.61% in accuracy, 81.91% in precision, and 91.05% in recall. Though the precision is 15% lower than the worst-performed deep learning model, its low training time cost and fast processing speed still shows its potential as part of the auxiliary indicator.

IV. FURTHER ANALYSIS AND DISCUSSION

Pancreas mis-segmentation can adversely affect classification results as it directly affects the ROI definition.

For instance, regions distant from the pancreas could be misclassified as pancreatic tissue. Such a situation causes an unnecessary increase in size of the ROI. Going one step further, this issue could lead to the inaccuracy of feature extraction in radiomics and the increasing time in training the deep learning model. This issue warrants further investigation for potential remedies. For example, a post-processing scheme can be applied. In our experiments, we did not encounter such effect.

Conversely, incomplete feature extraction from the pancreas and its surroundings on CT images results from the truncation issue. We empirically expanded the ROI area by 10% to address this issue. Despite exhibiting considerable performance in terms of precision, the model faced challenges in recall, displaying a bias toward classifying images with peri-pancreas edema as normal pancreas tissue. In general, the performance difference of the traditional CNN model is relatively large in accuracy and recall. What can be seen in the deep learning model's performance is the model with a simpler structure tends to reach higher scores. In particular, though transformer models consume more time in training, all the transformer models gain very decent scores in all metrics. If further development of the model is required, the model that performs well in this experiment will be a great reference in building the model's structure.

To further enhance the robustness of this experiment, it would be invaluable to explore variations in the model's performance by investigating different sizes for ROI expansion and diverse approaches to ROI extraction methodologies.

V. CONCLUSION

In this study, we established a very unique dataset comprising CT images from 255 pancreatic patients, categorized into classes indicative of the presence or absence of peri-pancreatic edema. Then, we employed LinTransUNet for the segmentation mask generation of CT pancreas images, yielding a commendable dice coefficient of 80.85%. These generated segmentation masks were then utilized for ROI extraction, which is used for constructing two types of peri-pancreatic edema prediction models. We evaluated different deep learning backbones in this study. The best deep learning model (Swin-tiny) demonstrates robust accuracy rates averaging 97.95%. While the radiomics-based model exhibited a lower score in every indicator, it achieved an accuracy of 79.61% and a precision of 81.91%. Furthermore, these segmentation masks were used to train the radiomics and deep learning-based classifier. By evaluating various deep learning and radiomics based models, we established a new baseline for peri-pancreatic edema detection.

REFERENCES

- [1] Koushik Biswas, Debesh Jha, Nikhil Kumar Tomar, Gorkem Durak, Alpay Medetalibeyoglu, Matthew Antalek, Yury Velichko, Daniela Ladner, Amir Bohrani, and Ulas Bagci. Adaptive smooth activation for improved disease diagnosis and organ segmentation from radiology scans. *arXiv preprint arXiv:2312.11480*, 2023.
- [2] Nicolas Carion, Francisco Massa, Gabriel Synnaeve, Nicolas Usunier, Alexander Kirillov, and Sergey Zagoruyko. End-to-end object detection with transformers. In *European conference on computer vision*, pages 213–229, 2020.
- [3] Tianqi Chen and Carlos Guestrin. Xgboost: A scalable tree boosting system. In *Proceedings of the 22nd acm sigkdd international conference on knowledge discovery and data mining*, pages 785–794, 2016.
- [4] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, et al. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929*, 2020.
- [5] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, Jakob Uszkoreit, and Neil Houlsby. An image is worth 16x16 words: Transformers for image recognition at scale. *ICLR*, 2021.
- [6] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016.
- [7] Andrew Howard, Mark Sandler, Grace Chu, Liang-Chieh Chen, Bo Chen, Mingxing Tan, Weijun Wang, Yukun Zhu, Ruoming Pang, Vijay Vasudevan, et al. Searching for mobilenetv3. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 1314–1324, 2019.
- [8] Gao Huang, Zhuang Liu, Laurens Van Der Maaten, and Kilian Q Weinberger. Densely connected convolutional networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4700–4708, 2017.
- [9] Debesh Jha, Nikhil Kumar Tomar, Vanshali Sharma, and Ulas Bagci. Transnetr: Transformer-based residual network for polyp segmentation with multi-center out-of-distribution testing. *arXiv preprint arXiv:2303.07428*, 2023.
- [10] Philippe Lambin, Ralph TH Leijenaar, Timo M Deist, Jurgen Peerlings, Evelyn EC De Jong, Janita Van Timmeren, Sebastian Sanduleanu, Ruben THM Larue, Aniek JG Even, Arthur Jochems, et al. Radiomics: the bridge between medical imaging and personalized medicine. *Nature reviews Clinical oncology*, 14(12):749–762, 2017.
- [11] Ze Liu, Han Hu, Yutong Lin, Zhuliang Yao, Zhenda Xie, Yixuan Wei, Jia Ning, Yue Cao, Zheng Zhang, Li Dong, et al. Swin transformer v2: Scaling up capacity and resolution. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 12009–12019, 2022.
- [12] Ze Liu, Yutong Lin, Yue Cao, Han Hu, Yixuan Wei, Zheng Zhang, Stephen Lin, and Baining Guo. Swin transformer: Hierarchical vision transformer using shifted windows. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 10012–10022, 2021.
- [13] Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun. Faster r-cnn: Towards real-time object detection with region proposal networks. In C. Cortes, N. Lawrence, D. Lee, M. Sugiyama, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 28. Curran Associates, Inc., 2015.
- [14] Mark Sandler, Andrew Howard, Menglong Zhu, Andrey Zhmoginov, and Liang-Chieh Chen. Mobilenetv2: Inverted residuals and linear bottlenecks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4510–4520, 2018.
- [15] Hugo Touvron, Thibaut Lavril, Gautier Izacard, Xavier Martinet, Marie-Anne Lachaux, Timothée Lacroix, Baptiste Rozière, Naman Goyal, Eric Hambro, Faisal Azhar, et al. Llama: Open and efficient foundation language models. *arXiv preprint arXiv:2302.13971*, 2023.
- [16] Lanhong Yao, Zheyuan Zhang, Ugur Demir, Elif Keles, Camila Vendrami, Emil Agarunov, Candice Bolan, Ivo Schoots, Marc Bruno, Rajesh Keswani, et al. Radiomics boosts deep learning model for ipmn classification. In *International Workshop on Machine Learning in Medical Imaging*, pages 134–143, 2023.
- [17] Lanhong Yao, Zheyuan Zhang, Elif Keles, Cemal Yazici, Temel Tirkes, and Ulas Bagci. A review of deep learning and radiomics approaches for pancreatic cancer diagnosis from medical imaging. *Current Opinion in Gastroenterology*, 39(5):436–447, 2023.
- [18] Zheyuan Zhang and Ulas Bagci. Dynamic linear transformer for 3d biomedical image segmentation. In *International Workshop on Machine Learning in Medical Imaging*, pages 171–180, 2022.
- [19] Zheyuan Zhang, Lanhong Yao, Elif Keles, Yury Velichko, and Ulas Bagci. Deep learning algorithms for pancreas segmentation from radiology scans: A review. *Advances in Clinical Radiology*, 2023.