Atlas-Guided U-Net++ with EfficientNetB5 for Automatic Pancreas Segmentation in Abdominal CT Scans

Felipe R. S. Teles[1], Neilson P. Ribeiro[1,2], Luana B. da Cruz[3], Geraldo B. Júnior[1], Anselmo C. de Paiva[1], João O. B. Diniz[1,2], Omar A. C. Cortes[2]

[1]Núcleo de Computacão Aplicada
Programa de Pós Graduação em Ciência da Computação
Universidade Federal do Maranhão (UFMA)

[2]Fábrica de Inovação
Instituto Federal do Maranhão (IFMA)

[3]Laboratório de Inteligência Computacional Aplicada (LICA)
Universidade Federal do Cariri (UFCA)

felipe.teles@nca.ufma.br

Abstract. *Pancreas segmentation in abdominal computed tomography images is challenging due to the organ's variability in shape, size, and position. This work proposes an automatic segmentation method based on a 2D Convolutional Neural Network (CNN) approach, consisting of three steps: (1) filtering non-pancreas slices using a CNN, (2) region of interest detection via a probabilistic atlas, and (3) final segmentation with U-Net++ with an EfficientNetB5 backbone. The method achieves a mean Dice coefficient of 78.59% and Recall of 79.12%, with a lower computational cost compared to 2.5 and 3D approaches. Thus, our results stand out among state-of-the-art methods, providing a computationally efficient and accurate solution for diagnosis and treatment planning.*

1. Introduction

According to the International Agency for Research on Cancer, pancreatic cancer is the sixth most lethal among registered cases [Bray et al. 2024]. In 2022, more than 450,000 deaths caused by the disease were reported [Daniel et al. 2024]. In this scenario, developing solutions that help specialists in diagnosis is essential, increasing the chances of successful treatment and patient survival.

Identifying the pancreas on medical images is a critical step in treatment, radiation therapy planning, and monitoring structural evolution [Silva et al. 2021]. Abdominal computed tomography (CT) is widely used in this process, as it provides detailed images that enable the differentiation of the pancreas from surrounding tissues, facilitating its delineation for more precise clinical analysis.

Deep learning-based models have been commonly adopted for many tasks in the medical field [Neto et al. 2024, Junior et al. 2024, Diniz et al. 2021]. In this context, Convolutional Neural Networks (CNNs) have proven to be effective, and continuous advancements have enabled the development of solutions to challenges such as automatic pancreas segmentation in CT scans. However, this task remains a significant challenge due to the high anatomical variability of the pancreas among patients, including differences in shape, size, and position. Furthermore, CNNs involve a high computational cost,

as feature extraction from three-dimensional images requires significantly more processing power than from two-dimensional images [Lee et al. 2021].

Thus, the main goal of this work is to propose a computational method that combines preprocessing and deep learning techniques for pancreatic segmentation. Our hypothesis is that by using scope reduction techniques, such as filtering out slices without the pancreas and employing a probabilistic atlas for initial localization of the pancreatic region, the performance of a segmentation network like U-Net++ can be enhanced. Thus, we believe that the proposed method offers the following contributions:

- Reduce unnecessary information by filtering out slices without the pancreas and employing a Probabilistic Atlas to refine the region of interest.
- A fully automatic 2D CNN-based pancreas segmentation method that outperforms most related work.

Furthermore, effective segmentation methods can enhance Computer-Aided Diagnosis (CAD) for clinicians, reduce operational costs for hospitals, and shorten waiting times for patients, eventually improving their quality of life.

2. Related Works

The pancreas segmentation task using deep learning primarily follows two principal approaches. The first considers the entire volumetric data as input, treating voxels as fundamental units for analysis. The second approach processes individual bi-dimensional slices, where pixels serve as the main features. This section presents the related works.

Li et al. [Li et al. 2021] proposed a hybrid method that integrates both atlas-based and neural network-based approaches. Their study used the NIH Pancreas-CT dataset, with binarization applied between $-100HU$ and $240HU$. Additionally, they evaluated the method on the Medical Segmentation Decathlon (MSD) dataset. The input consisted of 2.5D images, processed by an FCN network. The best Dice score achieved was 91.05%, with an average Dice of 71.36%.

Tian et al. [Tian et al. 2021] proposed a method based on Markov Chain Monte Carlo and CNNs to determine the location of the pancreatic region before segmentation. They employed a 3D U-Net and tested their method on the NIH Pancreas-CT dataset. In the preprocessing phase, they used superpixels to enhance feature extraction. Their approach achieved a Dice score of 78.13% in the test dataset.

Zhu et al. [Zhu et al. 2023] introduced an image quality control toolbox combined with a segmentation algorithm for pancreas. They used a 3D U-Net to segment the pancreas, evaluating their approach on both the MSD and NIH Pancreas datasets. The highest average Dice score obtained was 75.43% on the NIH dataset.

Kurnaz et al. [Kurnaz et al. 2024] proposed Pascal U-Net for pancreas segmentation. Unlike traditional U-Net, Pascal U-Net establishes connections through convolutional layers instead of direct links between the encoder and decoder. Their model was evaluated on the TCIA dataset, achieving an average Dice score of 71.35%.

Finally, Ferrara et al. [Ferrara et al. 2024] proposed using Deep Convolutional Neural Networks (DCNNs) as a preprocessing step to segment the pancreatic region

before identifying malignant tumors. They trained their model using the MSD dataset, adopting a 2.5D approach with an SRSNet architecture. The best Dice score was 66.8%.

While 3D models, such as those proposed by Tian et al. (2021), have demonstrated strong segmentation performance, they require substantial computational resources, making them less practical for large-scale applications. Other studies, including those by Li et al. (2021), Zhu et al. (2023), and Ferrara et al. (2024), explored 2.5D methods with neural networks, often incorporating preprocessing techniques like binarization or quality control. However, none of these works utilized a probabilistic atlas for localization or explicitly aimed at reducing computational complexity by filtering out unnecessary slices before segmentation.

In this context, our approach introduces an innovative combination of preprocessing and deep learning by integrating a probabilistic atlas for pancreas localization and employing filtering strategies to minimize computational burden. Unlike previous works, our method strategically balances segmentation accuracy and computational efficiency, addressing both challenges in the pancreas segmentation task effectively.

3. Materials and Method

To increase the efficiency of pancreas segmentation and reduce the computational cost, the techniques adopted consist of (1) dataset and preprocessing (2) reducing the number of slices through a classification CNN, (3) image crop through probabilistic atlas application to decrease the number of unnecessary features for the network (4) train a segmentation CNN and (5) evaluate the segmentation using metrics, Figure 1 illustrates the whole method proposed.

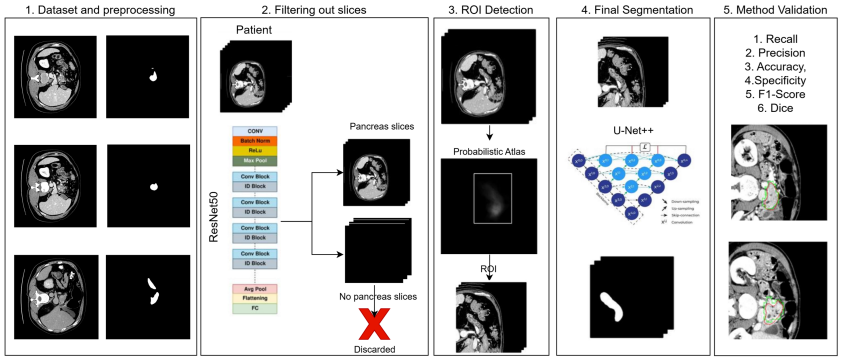

Figure 1. Illustration of method's step.

3.1. Dataset and Preprocessing

The dataset used in this study is the MSD [Antonelli et al. 2022], which contains 420 abdominal CT volumes collected by the Memorial Sloan Kettering Cancer Center. Each volume has approximately 100 slices (2d images) with 512×512 resolution. The pancreatic parenchyma was manually annotated in each slice by a radiologist using the Scout application. Examples of slices from the dataset can be seen in Figure 2, where the pancreas annotations vary between a single region or two distinct objects, depending on the slice.

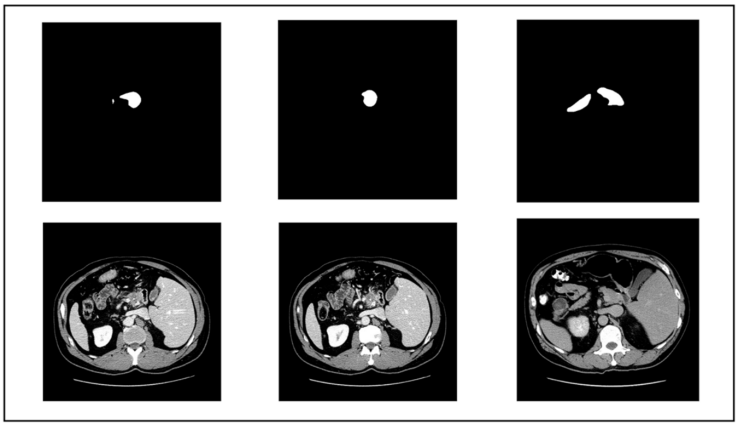

Figure 2. Examples from the dataset.

Before training the CNN models, the CT scans underwent a two-step preprocessing process. First, the intensity values of the volumes were limited to the range of $[-150, 250]$ on the Hounsfield Unit (HU) scale, improving the visualization of abdominal structures by excluding air regions and highly dense structures such as bones [Weston 2020, Fernandes et al. 2023]. In the second step, the CT scans were normalized to 8-bit values in the range of $[0, 255]$. In this study, CT slices were used in the proposed method.

3.2. Filtering out Slices

In this step, the slices per patient are reduced by filtering out those that do not include the pancreas, using a CNN trained for binary classification, where the positive class corresponds to a slice with the pancreas and the negative class to one without it. This step is crucial, as the majority of slices do not contain the organ, which could otherwise introduce bias and compromise the reliability of the final segmentation.

For this task, a ResNet50 model [He et al. 2016] was fine-tuned using pre-trained weights from ImageNet [Deng et al. 2009]. After training, the model was able to filter out slices that do not contain the pancreas, ensuring that only the relevant slices are preserved. This process not only improves computational efficiency by reducing unnecessary processing overhead but also contributes to a more sustainable and optimized segmentation pipeline.

3.3. ROI Detection

After selecting the slices that contain the pancreas, the Region of Interest (ROI) will be refined to cover only the area where the organ is located, optimizing the analysis scenario. To achieve this, we use a probabilistic atlas, a statistical technique that models the variability of a population, generating an average representation of specific characteristics.

By computing the average position of the masks in the training dataset, the atlas defines a probabilistic ROI for the pancreas, enabling the removal of irrelevant features through image cropping. The atlas is obtained using Equation 1:

$$A = \frac{1}{N} \sum_{i=1}^{N} \frac{I_i}{255} \tag{1}$$

where N is the total number of pixels in the area, and I_i denotes the intensity of the i-th pixel, normalized to the range $[0, 1]$.

A bounding box is then generated to define the ROI, with its dimensions constrained to multiples of 16 to maintain compatibility with the CNN input layer requirements. Finally, all filtered slices are cropped according to this ROI, preserving only the pancreatic region for the final segmentation step. Figure 3 illustrates the atlas with the corresponding bounding box.

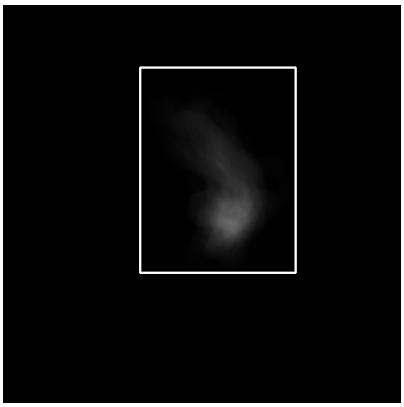

Figure 3. Probabilistic atlas representation.

3.4. Final Segmentation

After filtering out slices without the pancreas and detecting ROI using the probabilistic atlas, the next step is pancreas segmentation. For that, we employ U-Net++ [Zhou et al. 2018], an extension of the original U-Net [Ronneberger et al. 2015] combined with a Fully Convolutional Network (FCN) [Long et al. 2015]. Designed specifically for medical image segmentation, U-Net++ incorporates architectural refinements to enhance detail preservation and segmentation accuracy. The network architecture is illustrated in Figure 4.

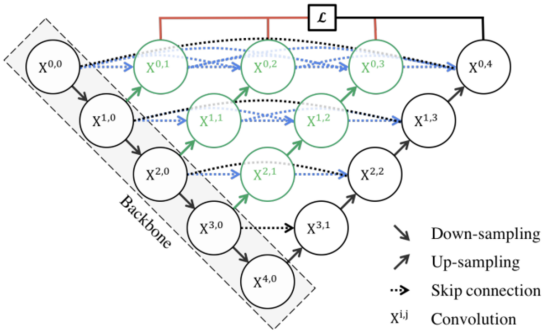

Figure 4. U-Net++ architecture [Zhou et al. 2018].

Traditional U-Net models may lose critical spatial information due to excessive down-sampling during training. To mitigate this, U-Net++ introduces dense skip pathways that enable a more gradual transition between encoder and decoder features. These pathways facilitate the aggregation of feature maps at different abstrac-

tion levels, improving segmentation accuracy, particularly for complex anatomical structures [Zhou et al. 2018].

An essential factor in U-Net++ is the selection of an appropriate backbone, which serves as the feature extractor and influences the model's capacity to learn hierarchical representations. Common choices include ResNet [He et al. 2016] and EfficientNet [Tan and Le 2019], both of which offer a balance between depth and efficiency. Results demonstrated that EfficientNetB5 provided the best performance for our method.

3.5. Method Validation

To validate the method, we used standard medical imaging metrics. In the slice filtering step, Accuracy, Specificity, Precision, Recall, and F1-Score assessed the model's performance in detecting slices with the pancreas while minimizing false classifications. For segmentation, the Dice Coefficient measured the overlap between the predicted mask and ground truth, while Recall ensured most of the pancreas was correctly segmented, reducing the risk of missing critical regions.

4. Results and Discussion

This section presents the results of the experiments to evaluate the proposed method. First, we define the training environment used for the development. Next, we present the experiments designed to validate each step. Finally, we compare our results with the related works discussed in Section 2 and provide case studies with a qualitative analysis.

4.1. Training Environment

The experiments were conducted using the Google Colab Pro cloud computing environment, which provides access to a Tesla T4 GPU with 12 GB of RAM. The implementation was written in Python, utilizing several key libraries: TensorFlow, PyTorch, NumPy, OpenCV, and Matplotlib.

4.2. Experiments

This section describes the different experimental scenarios designed to validate each step of the proposed method. Additionally, we analyze distinct backbones for U-Net++ to determine the most effective configuration for our approach.

To validate the method, we used a dataset described in Section 3.1 splited into 80% for training and 20% for testing, ensuring that slices from the same patient were not shared between sets, thus preventing data leakage. All techniques were applied only to the training set, while the test set was used exclusively for evaluation.

4.2.1. Experiment 1: Validation of the Filtering out Slices Step

In this stage, a ResNet50 was trained to classify whether the slices contained the pancreas or not (Section 3.2). The classification model was trained for 30 epochs using the Adam optimizer with a learning rate of 0.0001 and a batch size of 16. To find the optimal hyperparameters, we applied an optimization method based on Grid Search, which systematically explores a predefined set of values in the optimizer (Adam and AdamW), the learning rate (0.01, 0.001, and 0.0001) and the batch size (8, 12, and 16).

The ResNet50 architecture, initialized with ImageNet pre-trained weights, was subsequently fine-tuned on our target dataset. Figure 5 demonstrates consistent performance improvements across all evaluation metrics throughout the training epochs.

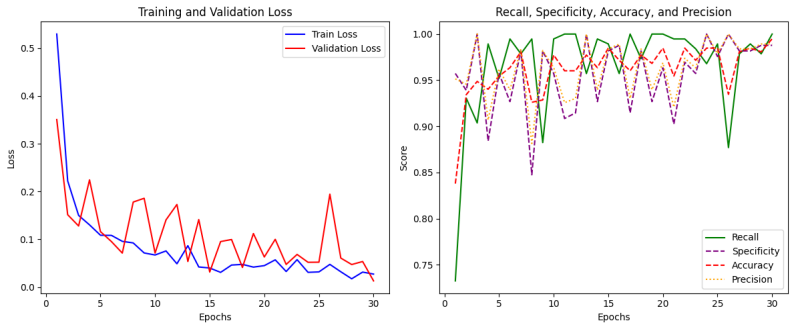

Figure 5. Result metrics per epoch.

Table 1 presents the performance metrics and their respective standard deviations across patients for the training and testing sets. In training, the model demonstrated a good balance between the metrics, indicating its ability to identify slices with the pancreas while minimizing false positives.

Table 1. Classification performance comparison in percentage (training vs. test).

Phase	Recall	Precision	F1-Score	Accuracy	Specificity
Training	96.03	98.81	99.22	97.30	95.57
Test	95.19 ± 0.01	96.22 ± 0.01	95.70 ± 0.01	95.44 ± 0.01	95.73 ± 0.02

In the testing set, the results were also consistent, with an F1-Score of 95.70%, showcasing the model's effective performance in distinguishing slices containing the pancreas. Additionally, the low standard deviation of the testing metrics indicates the stability and reliability of the model. These results highlight the importance of the filtering step, which contributes to reducing the number of slices to be processed in the segmentation, optimizing the efficiency of the method, and minimizing the presence of noise.

4.2.2. Experiment 2: Segmentation With and Without ROI Detection

After training ResNet50 to filter only pancreatic slices, the next step is to validate the effectiveness of the ROI Detection step (Section 3.3). To assess its impact, we conducted an experiment evaluating segmentation performance both with and without this step, using the U-Net++ proposed network.

For the comparison of experiments, we used the following configurations: a batch size of 16, AdamW optimizer, and a learning rate of 0.0003. In addition, learning rate optimization was performed using ReduceLROnPlateau to adjust it based on gradient dissipation. The results are summarized in Table 2, where best dice and best recall are the best patient result in test set.

Incorporating ROI Detection led to an additional performance boost, with 15% improvement in the mean Dice and Recall scores. This highlights the importance of the

Table 2. Segmentation performance in percentage with and without ROI detection.

Filtering Out	ROI Detection	CNN	Best Dice	Mean Dice	Best Recall	Mean Recall
Yes	No	U-Net++	73.35	62.32 $\pm$0.06	63.26	54.06 $\pm$0.07
Yes	Yes	U-Net++	83.86	78.59 $\pm$0.04	95.30	79.12 $\pm$0.07

ROI Detection step in filtering out irrelevant structures that could mislead the network. By focusing only on the pancreas region, the segmentation process becomes more precise, reducing interference from non-relevant anatomical structures and artifacts.

4.2.3. Experiment 3: Segmentation performance comparison U-Net++ vs. U-Net

In this second experiment, we evaluated the impact of U-Net++ with the standard U-Net architecture while keeping the Filtering Out and ROI Detection steps. This setup allows us to isolate the effect of using U-Net++ and determine whether it enhances the segmentation performance. U-Net configuration follows the same as U-Net++ (Section 4.2.2). The results are presented in Table 3.

Table 3. Segmentation performance in percentage comparison U-Net++ vs. U-Net.

Filtering Out	ROI Detection	CNN	Best Dice	Mean Dice	Best Recall	Mean Recall
Yes	Yes	U-Net	71.91	60.93 $\pm$0.05	86.11	64.03 $\pm$0.13
Yes	Yes	U-Net++	83.86	78.59 $\pm$0.04	95.30	79.12 $\pm$0.10

The results in Table 3 demonstrate that replacing the standard U-Net with U-Net++ significantly improves segmentation performance. The Dice coefficient increased from 60.93% to 78.59%, representing an improvement of approximately 17.66% in the mean Dice score. Similarly, the mean Recall improved by 15.09%, rising from 64.03% to 79.12%.

These findings confirm that U-Net++ enhances the segmentation quality by better capturing complex pancreatic structures. The improvements are likely due to its dense skip connections and deeper feature extraction capabilities, which enable more precise boundary delineation and better generalization across different patients. Consequently, incorporating U-Net++ into the proposed method is essential for achieving state-of-the-art segmentation accuracy.

4.2.4. Experiment 4: Evaluation of Different Backbone Architectures

In the final experiment, we compare different U-Net++ backbones to validate the choice of EfficientNetB5 for this task. Table 4 presents these results.

The results indicate that EfficientNetB5 achieved the best overall performance in both Dice and Recall metrics. It obtained the highest Best Dice score (83.86%) and Mean Dice (78.59%), demonstrating the best segmentation accuracy compared to the expert annotations. Also, it achieved the highest Recall (79.12%), indicating a greater accuracy in

Table 4. Segmentation comparison in percentage with different backbones.

Backbone	Best Dice	Mean Dice	Best Recall	Mean Recall
ResNet50	82.04	73.96 $\pm$0.07	93.01	76.49 $\pm$0.12
ResNet101	80.10	75.68 $\pm$0.07	96.65	78.64 $\pm$0.10
EfficientNetB3	82.99	71.49 $\pm$0.08	90.14	74.60 $\pm$0.12
EfficientNetB5	83.86	78.59 $\pm$0.04	95.30	79.12 $\pm$0.10

identifying the pancreas while generating fewer false positives. This suggests that EfficientNetB5 not only provides better alignment with expert annotations but also improves the reliability of segmentation by reducing errors. Given these findings, EfficientNetB5 was selected as the optimal backbone for this segmentation approach.

4.3. Comparison with Related Works

In this section, we compare the performance of our method with existing approaches in the literature in Table 5. The evaluation is based on the approach (3D, 2.5D, and 2D) and the Dice, highlighting the advantages and limitations of our approach.

Table 5. Comparison of dice scores for different approaches.

Authors	Dataset	Approach	Technique	Dice Score (%)
[Li et al. 2021]	NIH	2.5D	FCN	71.36
[Tian et al. 2021]	NIH	3D	3D U-Net	78.13
[Zhu et al. 2023]	MSD	3D	3D U-Net	75.43
[Kurnaz et al. 2024]	TCIA	2.5D	Pascal U-Net	66.80
[Ferrara et al. 2024]	MSD	2D	SRSNet	71.35
Proposed Method	**MSD**	**2D**	**Slice Filtering + ROI Detection + U-Net++**	**78.59**

Regarding the works described in Table 5, our proposed 2D method stands out as the best-performing approach, achieving a Dice score of 78.59%. This result surpasses both 2.5D and 3D methods, which are typically considered more robust architectures due to their ability to capture more complete spatial information. Moreover, our method offers a lower computational cost.

One of the key factors contributing to this performance is the effective of a slice filtering strategy, which removes images without the pancreas, reducing the presence of irrelevant data and improving training. Additionally, our method incorporates an Atlas-based approach for ROI detection, which localizes the pancreas more accurately before segmentation, minimizing false positives.

Finally, the backbone selection played a crucial role in our method's success. The use of EfficientNetB5 within the U-Net++ architecture provided a powerful feature extraction capability, enhancing segmentation accuracy without significantly increasing computational complexity. While direct comparisons are challenging due to differences in datasets, our results demonstrate the feasibility and effectiveness of a well-optimized 2D segmentation pipeline for pancreas segmentation.

4.4. Case Study

To further evaluate of the proposed method, we present case studies that provide a qualitative analysis of the segmentation results. These case studies illustrate how the method

performs in different scenarios and highlight its ability to accurately segment the pancreas in medical CT images. Figure 6 shows both the best and worst cases from the test dataset, where, in the slices, expert annotation is in red and segmentation is in green.

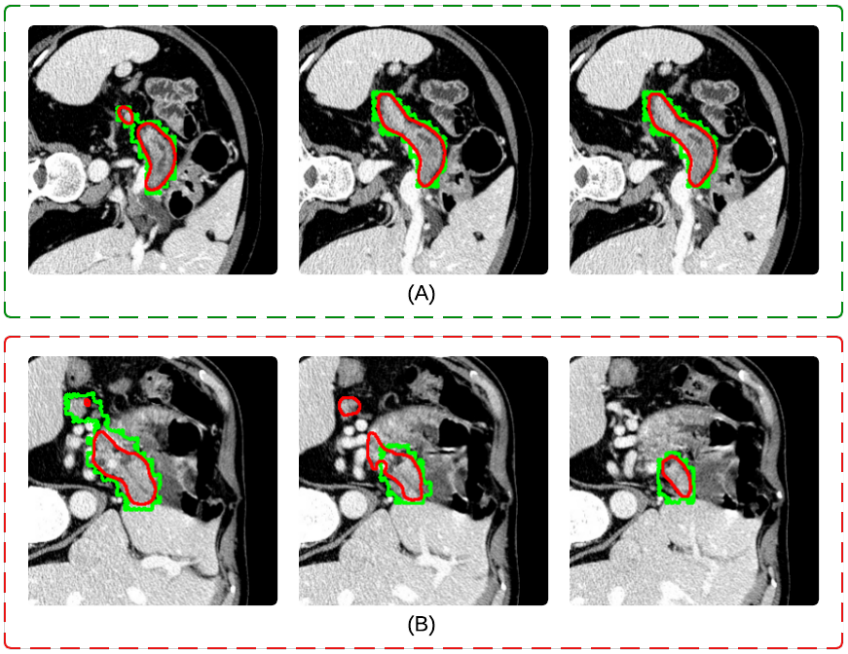

Figure 6. Case study: (A) Best result (B) Worst result.

In the example in Figure 6 (A), we present the best-case scenario, where the Dice score per slice is displayed. The method achieved a high degree of accuracy compared to the expert annotation, demonstrating its robustness. On the other hand, in Figure 6 (B), we observe a case where, despite successfully identifying an overlapping region with the annotation, the method over-segments the pancreas, marking a larger area. This issue occurred because the CT acquisition spacing caused significant overlap between the pancreas and surrounding organs with similar textures, making segmentation more challenging. A potential improvement to address this limitation could involve incorporating advanced techniques such as active contours to refine the segmentation process.

5. Conclusion and Future Work

Pancreas segmentation remains a challenging task due to the organ's irregular shape, low contrast, and anatomical variability. Additionally, the high computational cost associated with 3D models makes real-time clinical application difficult. The proposed method addresses these challenges by integrating slice filtering, probabilistic atlas-based ROI detection, and U-Net++ with an EfficientNetB5 backbone. This approach effectively reduces computational overhead while maintaining segmentation accuracy, achieving a best Dice score of 83.86% and a mean Dice score of 78.59%. Therefore, the proposed method is a viable alternative to computationally expensive 3D approaches and could be integrated into radiology workflows to assist in pancreas detection and treatment planning.

For future research, we propose evaluating the method on additional datasets. Additionally, incorporating post-processing techniques such as active contours or graph-based segmentation could refine the final mask. Exploring hybrid 2D/3D segmentation

strategies may further enhance performance by leveraging spatial consistency while preserving computational efficiency. These advancements could strengthen the clinical applicability of automated pancreas segmentation.

5.1. Acknowledgment

This work was carried out with the support of the Coordenação de Aperfeiçoamento de Pessoal de Nível Superior - Brasil (CAPES) - Código de Financiamento 001, Fundação de Amparo à Pesquisa do Maranhão (FAPEMA), Conselho Nacional de Desenvolvimento Científico e Tecnológico (CNPq).

References

Antonelli, M., Reinke, A., Bakas, S., Farahani, K., Kopp-Schneider, A., Landman, B. A., Litjens, G., Menze, B., Ronneberger, O., Summers, R. M., et al. (2022). The medical segmentation decathlon. *Nature communications*, 13(1):4128.

Bray, F., Laversanne, M., Sung, H., Ferlay, J., Siegel, R. L., Soerjomataram, I., and Jemal, A. (2024). Global cancer statistics 2022: Globocan estimates of incidence and mortality worldwide for 36 cancers in 185 countries. *CA: a cancer journal for clinicians*, 74(3):229–263.

Daniel, N., Farinella, R., Chatziioannou, A. C., Jenab, M., Mayén, A.-L., Rizzato, C., Belluomini, F., Canzian, F., Tavanti, A., Keski-Rahkonen, P., et al. (2024). Genetically predicted gut bacteria, circulating bacteria-associated metabolites and pancreatic ductal adenocarcinoma: a mendelian randomisation study. *Scientific Reports*, 14(1):25144.

Deng, J., Dong, W., Socher, R., Li, L.-J., Li, K., and Fei-Fei, L. (2009). Imagenet: A large-scale hierarchical image database. In *2009 IEEE Conference on Computer Vision and Pattern Recognition*, pages 248–255.

Diniz, J. O., Ferreira, J. L., da Silva, G. L., Quintanilha, D. B., Silva, A. C., and Paiva, A. (2021). Segmentação de coração em tomografias computadorizadas utilizando atlas probabilístico e redes neurais convolucionais. In *Anais do XXI Simpósio Brasileiro de Computação Aplicada à Saúde*, pages 83–94. SBC.

Fernandes, A. G. S., Braz Junior, G., Diniz, J. O. B., Silva, A. C., and Matos, C. E. F. (2023). Efficientdeeplab for automated trachea segmentation on medical images. In *Brazilian Conference on Intelligent Systems*, pages 154–166. Springer.

Ferrara, N., Andria, G., Scarpetta, M., Lanzolla, A. M. L., Attivissimo, F., Di Nisio, A., and Ramos, D. (2024). 2d and 2.5 d pancreas and tumor segmentation in heterogeneous ct images of pdac patients. In *2024 IEEE International Symposium on Medical Measurements and Applications (MeMeA)*, pages 1–5. IEEE.

He, K., Zhang, X., Ren, S., and Sun, J. (2016). Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778.

Junior, D. A. D., da Cruz, L. B., and Diniz, J. O. (2024). Classificação da camada lipídica do filme lacrimal usando k-means e deep learning. In *Simpósio Brasileiro de Computação Aplicada à Saúde (SBCAS)*, pages 1–12. SBC.

Kurnaz, E., Ceylan, R., Bozkurt, M. A., Cebeci, H., and Koplay, M. (2024). A novel deep learning model for pancreas segmentation: Pascal u-net. *Inteligencia Artificial*, 27(74):22–36.

Lee, H., Kim, Y. S., Kim, M., and Lee, Y. (2021). Low-cost network scheduling of 3d-cnn processing for embedded action recognition. *IEEE Access*, 9:83901–83912.

Li, J., Lin, X., Che, H., Li, H., and Qian, X. (2021). Pancreas segmentation with probabilistic map guided bi-directional recurrent unet. *Physics in Medicine & Biology*, 66(11):115010.

Long, J., Shelhamer, E., and Darrell, T. (2015). Fully convolutional networks for semantic segmentation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3431–3440.

Neto, C. M. S., Silva, A. L., Pessoa, A. C., Quintanilha, D. B., de Almeida, J. D., Junior, G. B., and Diniz, J. O. (2024). Diagnóstico de tuberculose em imagens de radiografia utilizando cvt. In *Simpósio Brasileiro de Computação Aplicada à Saúde (SBCAS)*, pages 342–353. SBC.

Ronneberger, O., Fischer, P., and Brox, T. (2015). U-net: Convolutional networks for biomedical image segmentation. In Navab, N., Hornegger, J., Wells, W. M., and Frangi, A. F., editors, *Medical Image Computing and Computer-Assisted Intervention – MICCAI 2015*, pages 234–241, Cham. Springer International Publishing.

Silva, G., Oliveira, F., Diniz, J., Diniz, P., Quintanilha, D., Silva, A., Paiva, A., and Cavalcanti, E. (2021). An automatic method for prostate segmentation on 3d mri scans using local phylogenetic indexes and xgboost. In *Anais do XXI Simpósio Brasileiro de Computação Aplicada à Saúde*, pages 165–176, Porto Alegre, RS, Brasil. SBC.

Tan, M. and Le, Q. (2019). Efficientnet: Rethinking model scaling for convolutional neural networks. In *International conference on machine learning*, pages 6105–6114. PMLR.

Tian, M., He, J., Yu, X., Cai, C., and Gao, Y. (2021). Mcmc guided cnn training and segmentation for pancreas extraction. *IEEE Access*, 9:90539–90554.

Weston, A. (2020). How to segment a pancreas ct. Accessed: 2025-02-11.

Zhou, Z., Rahman Siddiquee, M. M., Tajbakhsh, N., and Liang, J. (2018). Unet++: A nested u-net architecture for medical image segmentation. In *Deep Learning in Medical Image Analysis and Multimodal Learning for Clinical Decision Support: 4th International Workshop, DLMIA 2018, and 8th International Workshop, ML-CDS 2018, Held in Conjunction with MICCAI 2018, Granada, Spain, September 20, 2018, Proceedings 4*, pages 3–11. Springer.

Zhu, Y., Hu, P., Li, X., Tian, Y., Bai, X., Liang, T., and Li, J. (2023). An end-to-end data-adaptive pancreas segmentation system with an image quality control toolbox. *Journal of Healthcare Engineering*, 2023(1):3617318.