Volumetric Assessment of Pulmonary Artery Thrombus Burden

A Technical Report submitted to the Department of Biomedical Engineering

Presented to the Faculty of the School of Engineering and Applied Science

University of Virginia • Charlottesville, Virginia

In Partial Fulfillment of the Requirements of the Degree

Bachelor of Science, School of Engineering

Yeongju Lee

Spring, 2024

Technical Project Team Members

Rohey Colley

Nyota Patel

On my honor as a University Student, I have neither given nor received unauthorized aid on this assignment as defined by the Honor Guidelines for Thesis-Related Assignments

Timothy Allen, Department of Biomedical Engineering

Luke Wilkins, M.D., UVA Health Department of Interventional Radiology

Volumetric Assessment of Pulmonary Artery Thrombus Burden

Rohey Colley[1]. Yeongju Lee[1], Nyota Patel[1], Dr. Luke Wilkins[2], Dr. Kerstin Mueller[3]

[1] Department of Biomedical Engineering, University of Virginia
[2] Division of Vascular and Interventional Radiology, University of Virginia Hospital
[3] Siemens Healthineers

Abstract

Pulmonary embolism (PE) stands as a major cause of cardiovascular-related morbidity and mortality, affecting approximately 1-2 per 1,000 individuals annually in the U.S. alone [1]. The limitations of current thrombectomy procedures, including the inability to accurately quantify thrombus volume, significantly hinder effective treatment outcomes [1]. This project developed an artificial intelligence (AI) model designed to enhance the precision of thrombus volume quantification in PE using computed tomography (CT) imaging. By integrating advanced image preprocessing and deep learning-based segmentation using an adapted U-Net architecture, the model automates and refines the analysis of pulmonary artery thrombus burden. Our methods involved anonymizing CT scans from PACS, optimizing image quality through the skimage library in Python, and employing localized cropping to focus on areas of interest. The model was trained and evaluated using a comprehensive dataset consisting of variously preprocessed FUMPE images, with comparisons to manually segmented scans to validate accuracy. Preliminary results revealed that preprocessing significantly improved image clarity and contrast, while both localization and resizing increased the segmentation accuracy, evidenced by improved Intersection over Union (IoU) scores from 0.00026 to 0.00057 and from 0.00052 to 0.013, respectively. These enhancements contribute to a more reliable clinical assessment tool, potentially revolutionizing PE therapeutic strategies by providing clinicians with precise, quantifiable data to guide treatment decisions and improve patient outcomes.

Keywords: Pulmonary Embolism (PE), Artificial Intelligence (AI), Computed Tomography (CT), Image Segmentation, Thrombus Volume Quantification, Image Preprocessing, Intersection over Union (IoU)

Introduction

Pulmonary embolism (PE) remains a major public health concern due to its significant mortality and morbidity rates. It is predominantly caused by blood clots that travel to the lungs, obstructing pulmonary arteries and impairing gas exchange [Fig. 1]. The annual incidence of PE is about 60-70 cases per 100,000 individuals, with the condition accounting for nearly 100,000 deaths each year in the United States alone, making it the third leading cause of cardiovascular-related deaths [2].

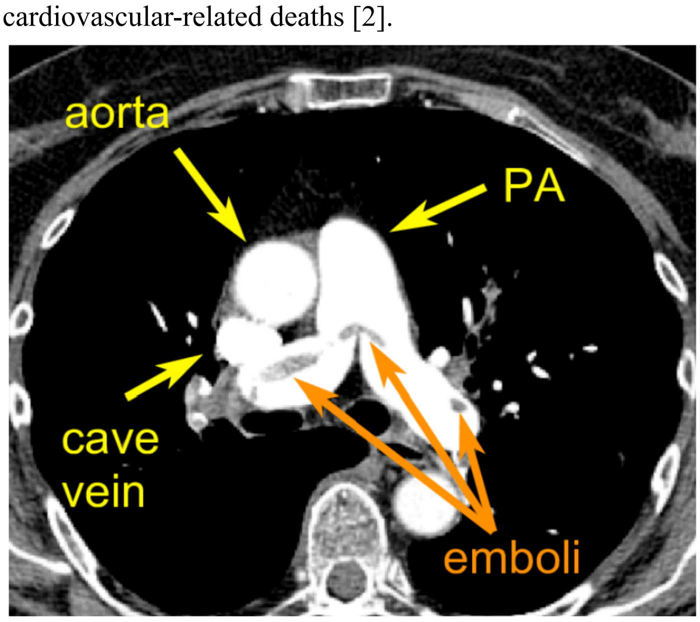

Fig. 1. CT image showing a pulmonary embolism in the pulmonary artery (PA) with labeled anatomical structures. The image clearly identifies the aorta, cave vein, and emboli, highlighting the locations relevant to the diagnosis of pulmonary embolism [3].

The traditional treatment modalities for PE include anticoagulation, thrombolytic therapy, and surgical or mechanical thrombectomy. Each method has its limitations. For instance, anticoagulants prevent new thrombi formation but do not dissolve existing clots, while thrombolytics carry a risk of significant bleeding. Mechanical thrombectomy has shown promise in rapidly removing clots and restoring hemodynamics but requires precise operation and has its own set of risks such as hemorrhage or damage to pulmonary vessels [1].

Recent advancements in mechanical thrombectomy, such as the introduction of devices like Inari Medical's FlowTriever, are noteworthy [Fig. 2]. This device facilitates the percutaneous removal of thrombi, offering a minimally invasive alternative to traditional methods. It has been associated with reduced procedure times and hospital stays, underlining the evolving landscape of PE management [4].

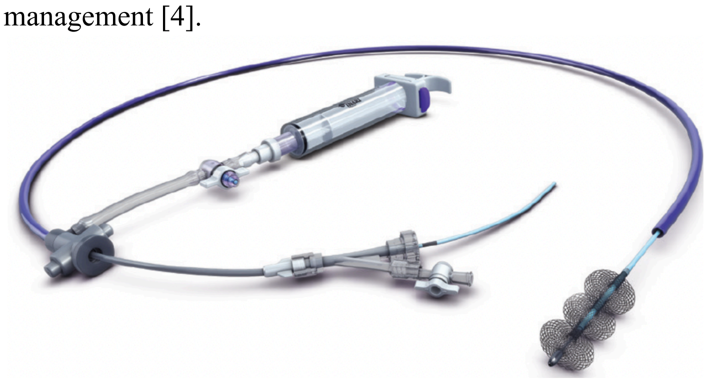

Fig. 2. Inari FlowTriever Device, designed for mechanical thrombectomy in pulmonary embolism treatment. This image shows the complete assembly, which includes a flexible catheter and a series of self-expanding nitinol discs designed to engage and remove thrombus efficiently from blood vessels.

Despite these advancements, a significant challenge persists in accurately quantifying the thrombus volume during and after interventions. Current practices often rely on subjective assessments or post-procedure imaging, which may not accurately reflect the outcome of thrombectomy. As Liu et al. (2022) point out, there is a notable absence of robust, automated methods to quantify thrombus volume on CT imaging datasets, leading to potential discrepancies in evaluating the efficacy of thrombectomy.

The literature reveals a growing interest in leveraging artificial intelligence (AI) to enhance imaging techniques. AI and machine learning (ML) models, particularly those involving convolutional neural networks (CNNs) like U-Net, have been effectively used for medical image segmentation. These models are trained to differentiate pathological features from normal anatomical structures, providing detailed and accurate analyses [5]. Furthermore, studies have explored the effectiveness of various image preprocessing techniques to enhance model accuracy, including methods like adjust_sigmoid for contrast enhancement and cropping strategies to focus on areas of interest, significantly impacting segmentation outcomes [5].

This project proposes to bridge the current technological gap by developing an automated AI-driven model capable of precise thrombus segmentation and volume quantification. By integrating deep learning techniques with existing CT imaging technology, the model aims to provide real-time, accurate assessments of thrombus burden. This advancement could potentially transform the therapeutic approach to PE, allowing for immediate and objective evaluation of thrombectomy success, optimizing patient management, and potentially improving clinical outcomes.

The need for precise, automated thrombus measurement is clear, given the clinical implications of accurate thrombus quantification on treatment strategies and patient outcomes. This project's integration of AI in medical imaging not only addresses a critical gap in pulmonary embolism management but also contributes to the broader field of medical diagnostics by enhancing the reliability and efficiency of imaging analyses.

Design Constraints/Assumptions/Limitations

In the development of an AI-driven solution for volumetric assessment of pulmonary artery thrombus burden, our project navigates through a series of design constraints, assumptions, and limitations that shape the scope and effectiveness of our outcomes.

Design Constraints

One of the primary constraints involves hardware and software compatibility. Our solution needs to integrate seamlessly with existing medical imaging systems, particularly Siemens CT scanners and picture archiving and communication system (PACS) . This requirement restricts our choice of development tools and demands adherence to medical industry standards, including data security and privacy protocols such as Health Insurance Portability and Accountability Act (HIPAA). Another significant constraint is data availability and quality. The project relies on a steady supply of high-quality, annotated medical imaging data, which is not consistently available. Variabilities in imaging protocols and scan qualities across different centers may impact the consistency and reliability of our AI model's training and performance. Moreover, the need for substantial computational resources, particularly GPUs for training deep learning models, poses a logistical and financial constraint that could influence project timelines and outcomes.

Assumptions

Our project hinges on several critical assumptions that significantly influence the design and expected outcomes. Firstly, we assume that the quality of the CT scans used for training and validation of the AI models is consistent and high across different datasets. This uniformity is crucial for ensuring that the AI model learns relevant features without biases introduced by variations in imaging quality. We also assume that the thrombi in pulmonary arteries present certain consistent characteristics in terms of density, shape, and location, which are recognizable by the AI algorithms. These characteristics are expected to be sufficiently distinct to enable accurate segmentation and volume measurement by the AI.

Additionally, there is an underlying assumption that the pre-processing techniques, such as normalization and enhancement of image contrast, adequately prepare the images for effective machine learning analysis without introducing artifacts that could mislead the AI model. We further assume that once trained, the model will perform with a similar level of accuracy and reliability in different clinical environments, a crucial factor for its broad application in medical diagnostics.

Limitations

The training data's potential bias represents a significant limitation. If the data does not encompass a wide spectrum of PE cases, including variations in patient demographics and thrombus characteristics, the model's performance may not be robust in real-world applications due to biases within the training dataset. The reliability of manual segmentation, used for creating ground truth labels, is susceptible to human error and inter-observer variability, introducing potential inaccuracies into the training data. Additionally, the complexity of deep learning models, such as the U-Net adapted for segmentation, risks overfitting to the training data, which could hamper the model's ability to generalize to new, unseen datasets. Scalability issues related to processing large imaging files and performing real-time analysis in clinical settings pose practical challenges. Lastly, the regulatory hurdles for medical software approval and the subsequent clinical adoption process are lengthy and complex, requiring rigorous validation to meet safety and efficacy standards.

Each of these elements—constraints, assumptions, and limitations—must be meticulously addressed to develop a robust, reliable, and clinically useful tool for assessing pulmonary artery thrombus burden. This careful management ensures that our project not only meets academic and clinical standards but also contributes meaningfully to improving the diagnosis and treatment of pulmonary embolism.

Material and Methods

In our project, we developed a sophisticated methodology to enhance the accuracy of thrombus volume quantification using computed tomography (CT) imaging for diagnosing PE. This involved a series of detailed steps starting from data collection, through image preprocessing, to the application of machine learning techniques.

Data Collection

The initial phase of our methodology focused on data collection and preparation, which is foundational for the accuracy of subsequent image processing and analysis. We sourced a variety of CT scans from the hospital's PACS system, ensuring a mix of 30 cases with and and 10 cases without PE to develop a robust model capable of distinguishing the condition effectively. Every scan was anonymized to comply with patient privacy protocols, an essential practice for upholding the ethical standards of our research. Additionally, to address a shortage of samples, we incorporated a public dataset known as FUMPE (Ferdowsi University of Mashhad's PE dataset). These scans were maintained in the Digital Imaging and Communications in Medicine (DICOM) format, a medical imaging standard that ensures compatibility with our processing tools.

Pre-Processing

The organization of DICOM data involved multiple folders, each containing several .dcm files representing individual scans. Our preprocessing script, developed in Python, was designed to iterate through each folder and file, sorting them by instance number to maintain the sequential integrity of the CT slices. This organization is critical for constructing accurate volumetric models from the slices.

Our preprocessing routine was instrumental in preparing the data for segmentation. We automated the loading of DICOM files from an input directory and implemented several transformations to optimize them for analysis. This included normalization of DICOM data to standardize the pixel intensity values across different scans, enhancing consistency in image analysis. We also enhanced image contrast using the **Skimage.exposure.adjust_sigmoid** function, which helped in making thrombi more distinguishable from the surrounding tissue [Fig. 3A, Fig. 3B]. Furthermore, we cropped the images to a fixed size of 256x256 pixels, centered on the most relevant part of the scan, and localized the analysis by extracting slices from the 25th to 75th percentile, focusing on areas most likely to contain thrombi [Fig. 3C]. These processed images were then saved as TIF files in a newly created output folder, ready for segmentation.

For thrombus segmentation, we adapted a U-Net neural network, which is particularly effective in medical image segmentation tasks. The network was customized to our specific needs to ensure precise identification of thrombi within the pulmonary arteries. We manually segmented CT scans using ITK-Snap to create a ground truth dataset, which was crucial for training and validating our AI model [Fig. 3D]. Our model underwent extensive training involving several phases: initially with 21,621 images of unlocalized FUMPE files from 28 patients, then refined with 2,294 localized images from the same patient set, and further specialized training with 346 images from positive PE cases focusing on localized and resized FUMPE files. Each training phase involved rigorous comparison of the model's predicted segmentations against manual segmentations to assess accuracy and make necessary adjustments.

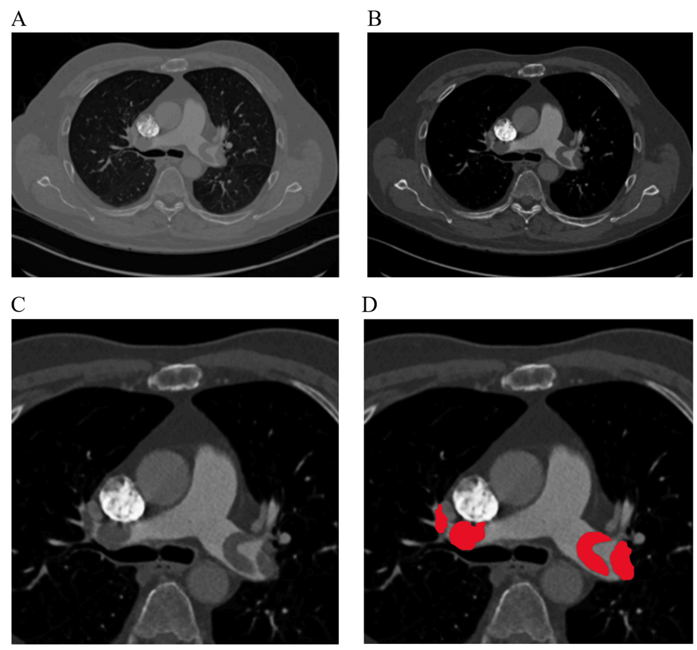

Fig. 3. Stages of CT imaging and processing for pulmonary embolism detection. A) Displays the original chest CT scan. B) Shows the scan after preprocessing enhancements to improve clarity and contrast. C) Illustrates the scan post-resizing for standardized analysis. D) Highlights the manually segmented thrombi (marked in red) using ITK-Snap, demonstrating targeted thrombus identification.

The final phase involved applying the trained model to new sets of CT scans to perform segmentation. The results were compared with manual segmentations to evaluate the model's effectiveness in real-world scenarios. This iterative process of training and testing helped refine the model to achieve high accuracy and reliability in segmenting pulmonary embolisms.

Results

The outcomes of our comprehensive testing and analytical processes have indicated significant advancements in both the methodological framework and practical application of our AI-driven approach in assessing thrombus in pulmonary embolism (PE) through computed tomography (CT) imaging.

Enhancement of Image Quality

Our project's pivotal focus was on enhancing the clarity and contrast of CT images to ensure the AI model could effectively distinguish thrombi from surrounding tissues. Among the various contrast enhancement techniques tested

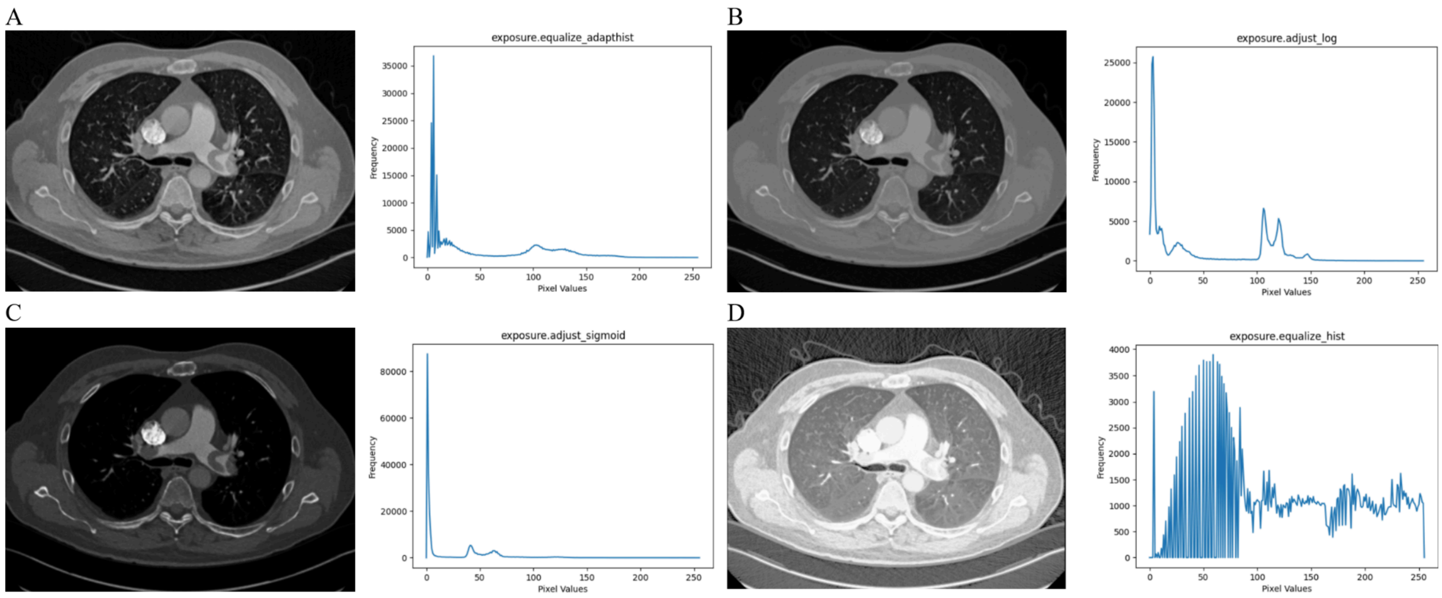

Fig. 4. Preprocessing techniques applied to a chest CT scan to enhance visibility of pulmonary emboli. A) Displays adapthist equalization with its pixel intensity histogram, subtly enhancing image contrast. B) Uses logarithmic adjustment (adjust_log) to highlight deeper structures. C) Shows sigmoid adjustment (adjust_sigmoid), emphasizing contrast between emboli and surrounding tissues. D) Presents histogram equalization (equalize_hist), with a histogram indicating uniform pixel value distribution, improving image clarity.

from skimage's exposure function—equalize_adapthist, adjust_log, adjust_sigmoid, and equalize_hist—adjust_sigmoid emerged as the most effective, significantly enhancing visual clarity [Fig. 4]. This improvement made thrombi more discernible, a critical factor for precise segmentation. Enhanced image quality leads to a reduction in model errors caused by ambiguous or poor-quality imaging data, thereby increasing the accuracy and reliability of our findings. This improvement in image quality is not only about enhancing visibility but also about ensuring that the enhanced images maintain a balance between contrast and detail retention without introducing noise or artifacts that could mislead the analysis. The visual benefits provided by this method are evident in Figure 3, which shows the distinct clarity and contrast improvements in processed images. Adjust_sigmoid's ability to enhance image detail also aids in identifying smaller or less dense thrombi that might be overlooked with less effective enhancement techniques. By achieving clearer images, our model's performance in detecting and segmenting thrombi is significantly improved, enhancing the overall effectiveness of thrombectomy evaluations.

Segmentation Accuracy Improvements

The targeted preprocessing localization strategy, which focused on the 25th to 75th percentile of slices, significantly improved the model's focus and relevance in imaging analysis. This strategic enhancement led to an increase in the Intersection over Union (IoU) scores from a mere 0.00026 to 0.00057, reflecting a substantial improvement in the model's ability to accurately identify and outline thrombus areas. These improved IoU scores are crucial for assessing the effectiveness of thrombectomy procedures, highlighting the precision of our AI model in clinical applications. The detailed improvements documented in Figure 5 exemplify how targeted preprocessing can mitigate issues related to irrelevant data and enhance the focus on critical imaging areas. The increase in segmentation accuracy is pivotal for clinicians who rely on precise data to make informed decisions about thrombectomy procedures. It also ensures that the AI system can be reliably used in clinical settings, where accurate segmentation directly influences treatment outcomes. Furthermore, this improvement reduces the likelihood of false positives and negatives, critical factors in medical imaging, especially in emergency medical situations like pulmonary embolism. By enhancing segmentation accuracy, we also facilitate more detailed volume calculations of thrombus burden, a direct benefit to determining the success of interventions. The ability to finely tune our AI model to recognize and accurately segment thrombi based on enhanced preprocessing illustrates our commitment to developing a reliable diagnostic tool.

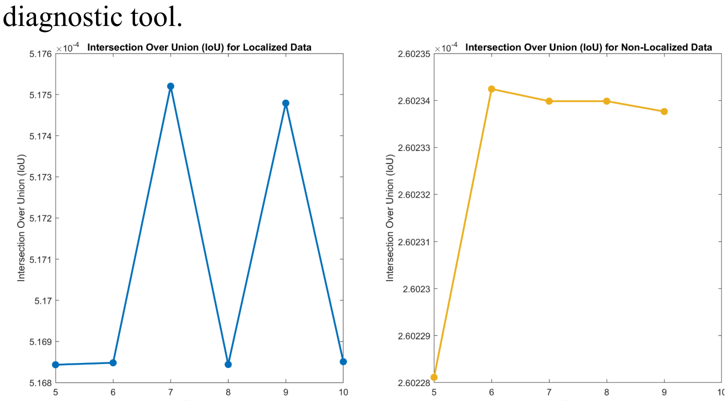

Fig. 5. Comparative IoU Trends for Localized and Non-Localized Data over Epochs. The left graph displays IoU scores for localized data, showing pronounced variability and a peak performance at the 7th epoch, suggesting optimal model learning at this stage. In contrast, the right graph for non-localized data demonstrates a stable, less variable performance trend. Note the significantly lower IoU scale in the non-localized data, indicating less precise segmentation. The absence of data for the 10th epoch in the non-localized series is due to memory constraints.

Influence of Resizing on Model Training and Testing

Cropping images to a central dimension of 256x256 pixels and then resizing them to a uniform 512x512 pixels during preprocessing created a more targeted and consistent dataset, which was instrumental in improving the training outcomes of our AI model. This methodological consistency led to an increase in IoU scores from 0.00051 to 0.012 during the training phase, showcasing the positive impact of standardization on the AI's learning capabilities [Tab. 1]. However, the improvement observed during training did not extend to the testing phase, where IoU scores remained at 0.00034, indicating a plateau. This suggests that while resizing aids in model training by providing uniform data, it does not necessarily enhance the model's ability to generalize to new, unseen datasets—a critical aspect for real-world application. The findings, as depicted in Table 1, underscore the need for further optimization of the model to ensure it performs consistently across different clinical scenarios. The discrepancy between training and testing phases also highlights potential overfitting during training, necessitating additional adjustments in the model's learning parameters. By addressing these challenges, we aim to enhance the model's robustness and reliability, ensuring it can effectively handle diverse clinical data. Additionally, this insight directs our future efforts towards refining preprocessing techniques and model parameters to better bridge the gap between laboratory settings and clinical applications. It also points to the importance of comprehensive validation studies to confirm the model's efficacy in accurately segmenting and analyzing thrombi in various patient populations.

Table 1. Comparison of Mean Intersection over Union (IoU) Scores for Training and Testing Phases. The table shows IoU scores comparing non-resized and resized image datasets during training and testing. The resizing process significantly enhances the IoU score during training, illustrating its positive impact on model learning, but shows no change during the testing phase, indicating a consistency in model performance across different datasets.

	Mean IoU Score	
Phase	**Non-Resize**	**Resize**
Train	0.0005176053	0.0125254997
Test	0.0003420587	0.0003420587

Volumetric Assessment and Future Directions

Our ongoing commitment to refine thrombus volume measurements through our AI model forms the foundation of our project, aiming to enhance the clinical assessment of thrombectomy procedures. Accurate quantification of thrombus volumes is essential for making informed clinical decisions and assessing the success of interventions. While our preprocessing and segmentation enhancements have significantly improved the foundation for these volumetric calculations, extensive studies and a broader array of CT scans are necessary to refine our techniques and validate their accuracy and reliability. The ultimate goal of integrating these assessments into clinical workflows is to provide a robust, real-time, precise measurement tool vital for effective management of PE. This aspect of our research is critical for moving forward and requires continuous development and refinement of our model, as well as an expansion of the training and testing datasets. These steps are crucial to further validate the model's accuracy and utility in clinical settings, ensuring it can make a significant impact on improving patient outcomes in pulmonary embolism management. The enhancements in volumetric assessment will allow clinicians to quantify thrombus burden more effectively, providing a measurable and repeatable method to evaluate the success of thrombectomy procedures. Our efforts to integrate these capabilities into standard clinical practice will facilitate a shift towards more data-driven, precise approaches in medical interventions, ultimately improving the safety and efficacy of treatments for pulmonary embolism. The detailed quantitative and qualitative results documented throughout our project, supported by figures, not only validate our methodological choices but also highlight the potential of our AI-driven approach to significantly advance PE treatment protocols.

Discussion

The capstone project undertaken by our team sought to address a critical challenge in the management of pulmonary embolism (PE)—the precise quantification of thrombus volume using computed tomography (CT) imaging. Our goal was to enhance the precision of catheter-directed thrombectomy procedures by providing a reliable computational framework for volumetric analysis. Over the course of this project, we have developed an advanced artificial intelligence (AI) model that leverages deep learning techniques for image segmentation combined with mathematical algorithms to calculate thrombus volumes accurately. This section concludes our findings and discusses the implications of our work, as well as future directions.

Technical Achievements

We successfully developed and implemented a series of preprocessing techniques that significantly enhanced the clarity and contrast of CT images, which are crucial for the effective segmentation of thrombi. The adjustment of image exposure and localization of the analysis to the most relevant slices were instrumental in improving the performance of our AI model. Our adapted U-Net neural network demonstrated robust capabilities in segmenting thrombus from CT scans, as evidenced by improved Intersection over Union (IoU) scores from initial testing phases.

Clinical Impact

The application of our AI model in clinical settings promises to revolutionize the treatment of PE by providing real-time, accurate assessments of thrombus burden. This capability allows for more informed decision-making during thrombectomy procedures and could potentially reduce the rates of recurrent PE by ensuring that sufficient thrombus material has been removed. Our system's ability to provide detailed volumetric data helps bridge the gap between the clinical assessment of thrombectomy success and actual patient outcomes, facilitating a more tailored approach to patient care.

Challenges and Learning

Throughout the project, we faced several challenges, particularly in data acquisition and the optimization of our AI model for diverse clinical scenarios. The scarcity of high-quality, annotated CT scans posed significant hurdles in training our model to the desired level of accuracy. However, these challenges also provided valuable learning opportunities, particularly in the realms of data handling and machine learning application in medical imaging. The iterative process of model training and validation helped refine our approach and improve the model's performance.

Future Work

Looking forward, there are several avenues for further research and development. First, expanding the dataset with more varied cases of PE will be crucial for enhancing the model's accuracy and generalizability. Additionally, integrating feedback mechanisms within the AI system could allow for continuous learning and improvement based on new patient data. Another key area for development is the automation of the segmentation process to reduce reliance on manual input, thereby decreasing potential biases and errors. Lastly, conducting comprehensive clinical trials to validate the effectiveness and reliability of our AI model in a real-world setting will be essential to its adoption in clinical practice.

Concluding Remarks

The implications of our work extend beyond the immediate benefits of improved thrombus measurement. By automating a critical aspect of PE management, our project contributes to the broader field of medical imaging and intervention, potentially setting new standards for the integration of AI technologies in healthcare. The success of this capstone project not only demonstrates the feasibility of advanced AI applications in a complex clinical domain but also highlights the transformative potential of such technologies in enhancing patient care and treatment outcomes. Our journey from concept to implementation reflects a significant achievement in biomedical engineering, paving the way for future innovations in the field.

End Matter

Author Contributions and Notes

R.C. performed image segmentation and wrote the paper.
Y.L. converted .dcm and .nrrd images of the CT scans into .tif images. Y.L also wrote the Python program to pre-process the CT scans, train and validate image segmentation using U-Net.
N.P. gathered, anonymized, and exported the CT scans as .dcm images. N.P. also helped with manual segmentation and wrote the Python program to load and pre-process the CT scans.

Acknowledgments

We extend our deepest appreciation to Dr. Timothy Allen, our BME Capstone course instructor, whose guidance and educational support were invaluable throughout the development of our project. We are equally grateful to Dr. Luke Wilkins, our faculty advisor, whose clinical insights into interventional radiology significantly enriched our understanding and application of the project objectives. Special thanks are due to Dr. Mueller from Siemens for her technical assistance and to our graduate teaching assistants, who provided crucial support and feedback. We also thank the University of Virginia's Biomedical Engineering department for providing the resources that supported our work. The dedication and assistance of everyone involved were instrumental to the success of our project.

References

[1] V. M. Chandra, M. S. Khaja, M. C. Kryger, A. K. Sista, L. R. Wilkins, J. F. Angle, and A. M. Sharma, "Mechanical aspiration thrombectomy for the treatment of pulmonary embolism: A systematic review and meta-analysis," *Vascular Medicine*, vol. 27, no. 6, pp. 574-584, 2022.

[2] A. K. Tarbox and M. Swaroop, "Pulmonary embolism," International Journal of Critical Illness and Injury Science, vol. 3, no. 1, pp. 69-72, 2013.

[3] Sebbe, R., et al., "Pulmonary arteries segmentation and feature extraction through slice marching," Proceedings of the 14th PRORISC Workshop on Circuits, Systems and Signal Processing, 2003.

[4] B. C. Wible, J. R. Buckley, K. H. Cho, M. C. Bunte, N. A. Saucier, and J. J. Borsa, "Safety and Efficacy of Acute Pulmonary Embolism Treated via Large-Bore Aspiration Mechanical Thrombectomy Using the Inari FlowTriever Device," *Journal of Vascular and Interventional Radiology*, vol. 30, no. 7, pp. 1073-1081, 2019.

[5] Z. Liu, H. Yuan, and H. Wang, "CAM-Wnet: An effective solution for accurate pulmonary embolism segmentation," *Medical Physics*, 2022.