

Automated Visceral Adipose Tissue Segmentation and Quantification from Abdominal MRI using an Enhanced U-Net and Region-growing

Sudha Devi B^{1*}, D. S Misbha²

¹ Department of Computer Science, Nesamony Memorial Christian College affiliated to Manonmaniam Sundaranar University, Abishekapatti, Tirunelveli 627 012, Tamil Nadu, India.
sudhanixen@gmail.com

² Department of Computer Applications, Nesamony Memorial Christian College affiliated to Manonmaniam Sundaranar University, Abishekapatti, Tirunelveli 627 012, Tamil Nadu, India.
misbhasatheesh4@gmail.com

KEYWORDS

VAT segmentation, U-Net, deep learning, attention mechanism, region growing, obesity, medical imaging.

ABSTRACT

The study focuses on enhancing the accuracy and reliability of visceral adipose tissue (VAT) segmentation and quantification from abdominal MRI images. Accurate segmentation of VAT is crucial for assessing obesity-related health risks, as traditional methods struggle with irregular shapes and varying intensities. The research utilizes a methodology consisting of three key modules: homomorphic filtering for intensity inhomogeneity correction, a U-Net architecture with attention mechanisms for primary segmentation, and a region-growing algorithm for refining segmentation. Homomorphic filtering effectively separates bias fields, enhancing image quality by transforming multiplicative artifacts into additive ones and removing them with high-pass filtering. This process ensures precise segmentation by maintaining high-frequency anatomical details. The U-Net model incorporates attention mechanisms and skip connections to focus on VAT regions, utilizing both local and global image contexts. The Combined Healthy Abdominal Organ Segmentation (CHAOS) challenge dataset and the Cancer Imaging Archive (TCIA) dataset are used to train and evaluate the model. It achieves a Dice Similarity Coefficient (DSC) of up to 0.985 on the CHAOS dataset and 0.972 on the TCIA dataset, outperforming existing methods in terms of segmentation accuracy. The region-growing algorithm further refines the segmentation by expanding VAT regions from high-confidence seed points, ensuring accurate boundary delineation and reducing noise. The study's results, evaluated using k-fold cross-validation, show that the proposed methodology significantly improves VAT segmentation efficiency, achieving a median DSC of 0.96 for the CHAOS dataset and 0.95 for the TCIA dataset in the most comprehensive experimental scenario. Comparative analysis indicates that the proposed approach outperforms other models, with higher sensitivity and specificity values, highlighting its potential for clinical applications in obesity management.

1. Introduction

Visceral adipose tissue (VAT), the fat that accumulates around internal organs, is recognized as a significant contributor to obesity-related health risks [1][2]. Unlike subcutaneous fat, which lies beneath the skin, VAT is strongly associated with various metabolic disorders, including cardiovascular disease, type 2 diabetes, and certain types of cancer [3][4]. As such, accurate assessment of VAT volume is crucial for understanding individual health risks and tailoring personalized treatment plans. Magnetic resonance imaging (MRI) has emerged as a valuable tool for non-invasive VAT quantification due to its ability to provide detailed anatomical information without exposing patients to ionizing radiation [5]. However, analyzing MR images for VAT segmentation and quantification is not possible without its challenges. The inherent variability in VAT distribution across individuals, coupled with image artifacts like intensity inhomogeneity, can impede accurate identification and measurement [6][7]. As illustrated in Figure 1, abdominal MRI scans of individuals with obesity often exhibit complex VAT morphology, highlighting the challenges in accurate segmentation and quantification [8]. Additionally, traditional manual segmentation techniques are time-consuming and prone to human error, highlighting the need for automated and reliable solutions [9]. Previous research efforts have explored various approaches to automate VAT segmentation in MR images. These include classical image processing techniques, such as thresholding and region growing, as well as more recent advances in machine

learning, particularly deep learning models like U-Net [10][11][12]. While these methods have shown promise, they often face limitations in accurately delineating VAT boundaries, especially in cases with complex shapes, varying intensities, and overlaps with other tissues [14][15]. Moreover, the presence of intensity inhomogeneity in MR images can further degrade the performance of these methods.




Figure 1 (a): Abdominal MRI of an Individual with Obesity

Figure 1(b): Illustration of the Complex Morphology of VAT in Abdominal MRI (Blue Overlay)

The proposed methodology offers a robust solution to the challenges of VAT segmentation and quantification in abdominal MR images. The integration of three key modules—bias field correction, U-Net segmentation, and VAT quantification—enables a precise and reliable assessment of VAT. Intensity inhomogeneity, a common artifact in MR images, can significantly hinder accurate VAT delineation. The proposed method addresses this issue by employing homomorphic filtering for bias field correction. This technique effectively removes the smooth, spatially varying intensity variations that can obscure VAT boundaries, thereby improving the quality of the input data for subsequent segmentation. The core of the methodology lies in the optimized U-Net architecture, designed specifically for VAT segmentation. Unlike traditional U-Nets, this model incorporates multiple attention blocks at various stages, enabling it to focus on subtle VAT patterns amidst complex anatomical structures. The attention mechanism enhances the model's sensitivity to the subtle features that distinguish VAT from surrounding tissues. Additionally, skip connections preserve fine-grained details during downsampling, ensuring sharp and accurate VAT boundaries in the final segmentation.

Following the U-Net segmentation, the region-growing algorithm further refines the VAT boundaries. This refinement step is crucial for addressing potential inaccuracies in the initial segmentation, such as over- or under-segmentation. By iteratively expanding the VAT regions based on similarity criteria, the algorithm ensures a more precise delineation of the VAT, leading to accurate volume quantification. The final module focuses on quantifying the segmented VAT volume. It converts the segmented regions into a binary mask, where each pixel represents either VAT or non-VAT tissue. The volume of each voxel is then calculated based on the image resolution, and the total VAT volume is obtained by summing the volumes of all VAT voxels. This quantitative assessment provides clinicians with a valuable tool for diagnosing and monitoring obesity-related health risks.

According to the results, the proposed method excels in its ability to deliver high segmentation accuracy and efficiency across different datasets. It achieves optimal performance within 300 epochs, indicating a well-tuned learning process. The method consistently outperforms existing approaches, particularly in complex cases involving VAT and non-VAT regions. The incorporation of k-fold cross-validation further strengthens the model's generalization ability, as seen in the consistent improvement of performance metrics with increasing folds. The combination of preprocessing with homomorphic filtering and region-growing significantly enhances the segmentation results, reducing errors in boundary delineation and increasing sensitivity and specificity. Additionally, the proposed method's superior DSC and Intersection over Union (IoU) highlight its ability to accurately capture VAT regions, outperforming other state-of-the-art models.

This paper is organized as follows: Section 2 reviews related work, highlighting the strengths and limitations of existing methods for VAT segmentation. Section 3 presents the details of the proposed methodology, including the integration of bias field correction, the U-Net architecture, and the region-growing algorithm. Section 4 outlines the experimental setup and discusses the results of the proposed method in

comparison with other approaches. Section 5 concludes the paper by summarizing the contributions and suggesting potential directions for future work.

2. Literature review

Hao Shen et al. [15] developed a deep-learning model for the automatic segmentation of abdominal muscle and fat. The developed algorithm, SECANet, is based on the U-Net architecture and incorporates a selective efficient channel attention block to improve segmentation accuracy. The methodology involves preprocessing the abdominal CT and MR images by adjusting Hounsfield unit values and applying z-score normalization. The main drawback of this model is its difficulty in accurately segmenting VAT regions, particularly in cases with smaller VAT areas. This issue occurs because the U-Net model struggles to handle unbalanced pixel proportions in the dataset, resulting in poor VAT segmentation. Additionally, the model tends to over- or under-segment complex regions due to insufficient feature extraction during training on diverse anatomical variations, which leads to errors in delineating VAT boundaries.

Prakash et al. [16] introduced a MultiRes-Attention U-Net to automate the segmentation of abdominal fat compartments, particularly superficial subcutaneous adipose tissue (SSAT), deep subcutaneous adipose tissue (DSAT), and VAT using MRI scans. The method integrates MultiRes blocks, attention gates, and a hybrid loss function to handle varying shapes and sizes of fat deposits across abdominal regions. The method demonstrated good generalization, but one key drawback is its missing segmentation in complicated regions. This issue arises because the MultiRes-Attention U-Net, despite using attention gates, struggles with intricate anatomical boundaries and the discontinuous nature of VAT, particularly around internal organs. Additionally, the model's dependence on 2D slice conversion from 3D MR scans reduces the ability to capture complex spatial relationships, leading to segmentation inaccuracies in challenging regions.

Kafali et al. [17] developed an automated method for abdominal VAT segmentation and volume quantification using MRI data. They proposed two 3D CNN models: ACD 3D U-Net and 3D nnU-Net. These models utilize volumetric multi-contrast MRI inputs and attention mechanisms to segment both subcutaneous and VAT. However, the authors did not fully address the complexity of anatomical structures in MRI scans. The U-Net models face challenges with intensity variations and overlaps, leading to inaccuracies in VAT boundary delineation.

Somasundaram et al. [18] developed an automated method for body composition analysis using nnU-Net for segmentation of VAT, subcutaneous adipose tissue (SAT), liver, and muscles based on water-fat MRI. The method utilized chemical shift-encoded MRI to extract quantitative metrics such as proton density fat fraction (PDFF) and organ volumes. However, the proposed method struggles with complex VAT boundaries, which often lead to over- or under-segmentation. This issue arises because the algorithm primarily relies on water-fat separation, neglecting complex anatomical structures and intensity inhomogeneities.

Ogunleye et al [19] proposed an automated approach for VAT and SAT quantification in adolescents using Dixon-based MRI and CNN. The method integrates region- and pixel-based segmentation, utilizing a U-Net architecture to segment abdominal regions and classify adipose tissue. A major drawback of the approach is its focus on single MRI slices, which limit the algorithm's ability to capture complex VAT boundaries.

Wu et al. [20] propose an automated deep-learning method to segment SAT and VAT from Dixon MRI scans in adolescents. The model, based on the 2D-CDFNet architecture, focuses on abdominal fat segmentation using axial fat and water images. However, the primary drawback relates to the complexity of VAT segmentation in adolescents, which presents challenges for precise boundary delineation. The segmentation quality is further compromised by the lack of consideration for anatomical variations that occur during puberty.

Christine Haugen et al. [21] method aims to segment and quantify adipose tissue (AT) in both mice and human MRI images. It utilizes a semi-automated approach involving image thresholding, background noise removal, object labelling, and manual editing to distinguish SAT from VAT. The algorithm calculates the volume of each AT depot by counting the voxels after segmentation. However, a major drawback is the reliance on manual editing, which introduces subjectivity and potential variability between users. This can lead to inconsistencies in the results.

Ina Vernikouskaya et al. [22] present a method for automated segmentation and quantification of abdominal body fat compartments, differentiating SAT and VAT, from T1-weighted MRI using a U-Net. The study successfully applied the algorithm to a dataset of both healthy controls and patients with amyotrophic lateral sclerosis (ALS), demonstrating a significantly increased VAT/SAT ratio in the ALS group. However, a potential drawback is the model's reliance on high-quality reference segmentations for training and validation. The limitations of these reference segmentations, such as those caused by motion artifacts or poor image contrast, could potentially impact the accuracy of the model's predictions.

Hwang et al. [23] developed a method for VAT segmentation using 3D CT images. The method aims to provide a VAT volume measurement by automatically dividing three anatomical compartments: the lungs, soft tissues, and post-vertebral spaces. The approach involves a three-step process, utilizing 3D CT images to separate the regions. However, this method faces segmentation issues, particularly in regions where VAT and SAT meet. These errors arise because the algorithm relies primarily on Hounsfield unit (HU) values, which do not effectively calculate the complex boundaries between different tissue types.

Several deep learning models and algorithms have been developed to address VAT segmentation challenges in abdominal scans from MRI and CT images. However, these approaches often face difficulties in handling complex anatomical boundaries, intensity variations, and the discontinuous nature of VAT. Despite their advancements, issues like over and under-segmentation in complex regions persist, highlighting the need for more refined techniques.

3. Proposed Method

The proposed methodology consists of three main modules. The initial module addresses intensity inhomogeneities in MR images using homomorphic filtering. Separating the bias field from anatomical details enhances image quality and ensures accurate subsequent analysis. An optimized U-Net architecture, incorporating attention mechanisms and skip connections, performs the primary VAT segmentation.

This deep learning model utilizes both local and global image features to identify VAT regions accurately, even in challenging cases of irregular shapes and varying intensities. The final module quantifies the segmented VAT volume. It converts the segmented regions into a binary mask, calculates the volume of individual voxels based on image resolution, and sums these volumes to obtain the total VAT volume. This provides a clinically relevant metric for assessing obesity-related health risks. Figure 2 shows the overall process flow of the proposed automated VAT segmentation and quantification system.

3.1 Dataset details

The dataset for this research includes two well-known abdominal MRI datasets. From the dataset, obesity patient MRI scans were chosen with the assistance of a skilled radiologist. The first dataset is from the Combined Healthy Abdominal Organ Segmentation (CHAOS) challenge [27]. This dataset contains abdominal CT and MRI (T1 and T2 weighted) images, but only the MR images are used in this study for model training and testing. The images were acquired using a 1.5T Philips MRI scanner, producing 12-bit DICOM images with a resolution of 256 x 256 pixels. The slice thickness varies between 5.5 and 9 mm, with an average of 7.84 mm. The x-y pixel spacing ranges from 1.36 to 1.89 mm, averaging 1.61 mm. The dataset includes between 26 and 50 slices per sequence, with an average of 36 slices. The second dataset is obtained from the Cancer Imaging Archive (TCIA), which offers a wide variety of abdominal MRI scans [28]. The second dataset was sourced from The Cancer Imaging Archive (TCIA), a publicly accessible repository of medical images. The combination of these two datasets enhances the model's ability to generalize to unseen data, ensuring its robustness and applicability in real-world scenarios.

3.2 Bias Field Correction Using Homomorphic Filtering

Bias field correction is crucial for subsequent VAT segmentation and quantification, ensuring that the resulting measurements are accurate and reliable. The bias field, also known as intensity inhomogeneity, is a low-frequency artifact that causes intensity variations across the image, leading to inaccurate segmentation and quantification. In this research homomorphic filtering is used to separate the multiplicative bias field from the true image, effectively correcting these intensity variations. Homomorphic filtering operates by transforming the multiplicative components of the image into additive ones, applying a high-pass filter to

remove the bias field, and then transforming the components back to their original domain [25]. This process ensures that the bias field, which varies slowly across the image, is separated from the high-frequency anatomical details. Convert the multiplicative bias field problem into an additive one using a logarithmic transformation. Let $I(x, y)$ be the observed image, $R(x, y)$ be the true image, and $B(x, y)$ be the bias field:

$$I(x, y) = R(x, y) \cdot B(x, y) \quad (1)$$

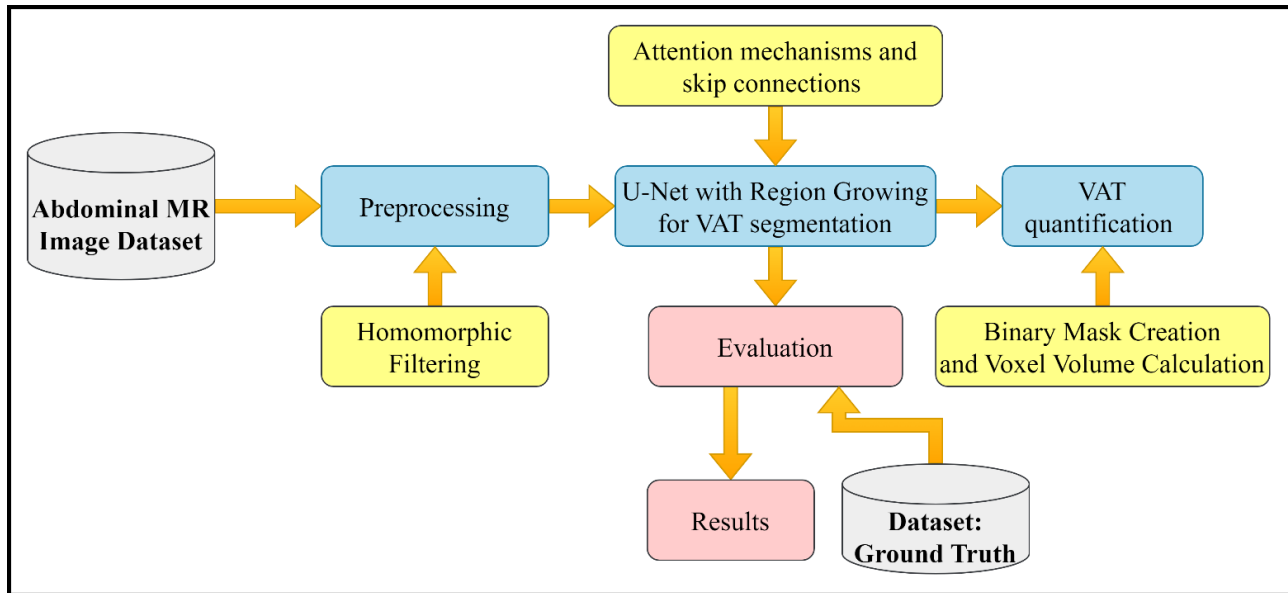


Figure 2. Overall process flow of proposed Automated VAT Segmentation and Quantification from Abdominal MRI using an Enhanced U-Net and Region Growing Algorithm

Applying the logarithm:

$$\log(I(x, y)) = \log(R(x, y)) + \log(B(x, y)) \quad (2)$$

Transform the logarithmically converted image into the frequency domain using the Fourier transform:

$$F(u, v) = \mathcal{F}\{\log(I(x, y))\} \quad (3)$$

Apply a high-pass filter to $F(u, v)$ to attenuate the low-frequency components (bias field) while preserving the high-frequency components (true image details). Let $H(u, v)$ be the high-pass filter:

$$F_H(u, v) = F(u, v) \cdot H(u, v) \quad (4)$$

Transform the filtered image back to the spatial domain using the inverse Fourier transform:

$$\log(I_H(x, y)) = \mathcal{F}^{-1}\{F_H(u, v)\} \quad (5)$$

Apply the exponential function to revert the logarithmic transformation, resulting in the corrected image:

$$I_H(x, y) = \exp(\log(I_H(x, y))) \quad (6)$$

By applying homomorphic filtering, the bias field correction module effectively removes intensity inhomogeneities from the MRI images. Homomorphic filtering enhances the overall quality of the medical images, making them suitable for precise analysis and diagnosis. This process addresses the challenge of irregular shapes, overlaps, and variability in VAT, providing a consistent and robust approach to bias field correction in medical imaging. Figure 3 shows the abdominal MR image before and after pre-processing.



Figure 3. (A) Original abdominal MR image before pre-processing; (B) noise-removed abdominal MR image (after bias field correction).

3.3 Proposed U-Net architecture

The proposed U-Net architecture for VAT segmentation from MR images builds upon the established U-Net framework, introducing key modifications to enhance accuracy and efficiency in handling the unique challenges of VAT detection [26]. The input layer accepts 320x320 pixel medical images with two channels. The initial convolutional block applies filters to extract essential features from the input data, transforming the dimensions to 320x320x128. This block not only captures low-level features like edges and textures but also introduces an attention mechanism, allowing the model to focus on regions of interest within the image.

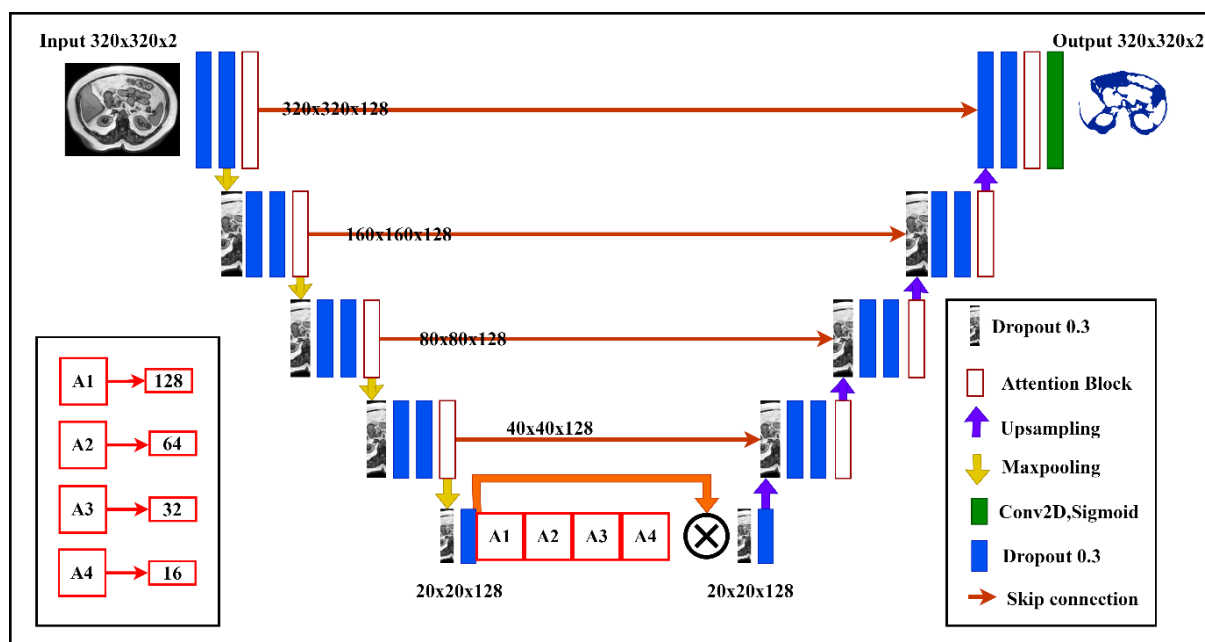


Figure 4 (a) Proposed attention-based U-net architecture for VAT segmentation from abdominal MR images.

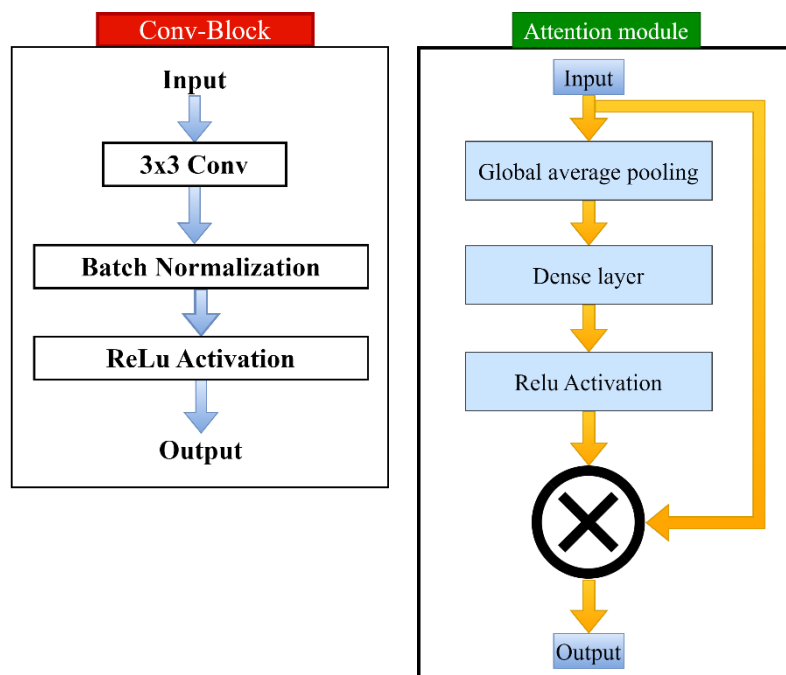


Figure 4 (b) Convolution block details of proposed U-net architecture.

Figure 4 (c) Attention block details of proposed U-net architecture.

The downsampling path progressively reduces the spatial dimensions of the feature maps through max-pooling, maintaining a consistent depth of 128 channels. This hierarchical approach allows the model to capture both local details and broader contextual information. The midblock, a crucial component, incorporates multiple attention blocks (ABs) with varying channel depths. Each AB utilizes global average pooling to summarize features, followed by dense layers and activation functions to generate attention maps. These maps modulate the input features, emphasizing relevant information and suppressing noise. The multiple ABs with different channel depths enable the model to capture attention at various scales, contributing to a more detailed understanding of VAT distribution. The upsampling path mirrors the downsampling path, increasing the spatial dimensions while refining the segmentation details. Dropout layers help mitigate overfitting by randomly deactivating neurons during training. Skip connections facilitate the fusion of high-level semantic information from the midblock with low-level spatial details from the downsampling path, ensuring accurate boundary delineation. The final output layer generates a segmentation map with the same dimensions as the input image. The sigmoid activation function ensures pixel values represent probabilities belonging to the VAT class, enabling precise localization of VAT regions.

The proposed U-Net architecture for VAT segmentation from MR images has several key enhancements. Figure 4 (a) shows the proposed attention-based U-net architecture for VAT segmentation from abdominal MR images. Figure 4 (b) shows the convolution block details of the proposed U-net architecture and Figure 4 (c) shows the attention block details of the proposed U-net architecture.

Multiple attention blocks integrated at various stages improve the model's ability to detect subtle VAT patterns amidst complex intensity variations often seen in MR images. Using two-channel input data allows for incorporating additional information from MR images, potentially enhancing tissue contrast and segmentation accuracy. Skip connections preserve fine-grained details during downsampling, ensuring sharp and accurate VAT boundaries in the final segmentation. An optimized midblock with varying attention block depths enables efficient capture of contextual information at different scales, facilitating the handling of diverse VAT shapes and sizes. These combined improvements result in a robust and accurate solution for VAT segmentation from MR images.

3.4 Region Growing algorithm and post-processing

The region-growing module in the proposed methodology focuses on refining the initial segmentation of VAT obtained from the U-Net model. This module addresses the unique challenges posed by irregular shapes, overlaps, and variability in VAT regions, ensuring that the final segmentation is both accurate and clinically useful. Region growing is a pixel-based image segmentation technique that starts with an initial set of seed points and expands these regions by including neighboring pixels that meet a certain similarity criterion. This method is particularly effective for medical imaging, where precise delineation of tissue boundaries is crucial.

Seed Point Selection: The process begins with selecting seed points within the VAT regions identified by the U-Net model. These seed points serve as the starting locations for the region growing process. In the context of VAT segmentation, seed points can be chosen based on high-confidence areas within the U-Net output, where the probability of being VAT is significantly high.

Similarity Criterion: A similarity criterion determines whether neighboring pixels should be included in the growing region. For VAT segmentation, this criterion is typically based on intensity values and spatial proximity. Let $I(x,y)$ denote the intensity of a pixel at location (x,y) , and I_{seed} be the intensity of the seed pixel. A pixel (x',y') is included in the region if it satisfies:

$$|I(x',y') - I_{seed}| \leq \epsilon \quad (7)$$

where ϵ is a predefined threshold. This ensures that only pixels with similar intensity values to the seed point are included in the VAT region.

Region Growing Algorithm: The algorithm iterates over the image, starting from the seed points and expanding the regions by including neighboring pixels that meet the similarity criterion. The process continues until no more pixels can be added to any of the regions.

Post-Processing: After the region growing algorithm completes, the resulting binary mask M may contain small isolated regions or noise. A morphological opening operation (erosion followed by dilation) can be applied to remove these artifacts and ensure that the VAT regions are smooth and well-defined.

Algorithm 1 Region Growing Algorithm

```

Set the initial seed points  $S = \{(x_i, y_i)\}$  based on the U-Net output.
2: Create an empty binary mask  $M$  of the same dimensions as the input image, initialized to zero.
3: Define a queue  $Q$  to manage the region growing process.
4: for each seed point  $(x_i, y_i) \in S$  do 5:   Add  $(x_i, y_i)$  to  $Q$ .
6: end for
7: while  $Q$  is not empty do
8:   Remove the front element  $(x, y)$  from  $Q$ .
9:   for each neighbor  $(x', y')$  of  $(x, y)$  do
10: if  $(x', y')$  is within image bounds and  $M(x', y') = 0$  and  $|I(x', y') - I(x, y)| \leq \epsilon$  then
11:   Set  $M(x', y') = 1$ .
12:   Add  $(x', y')$  to  $Q$ .
13: end if
14: end for
15: end while

```

By incorporating region growing, the proposed methodology ensures that VAT segmentation is refined and accurate, even in challenging cases with irregular shapes and varying intensities. This module complements the initial U-Net segmentation by addressing local intensity variations and spatial continuity, leading to more precise and clinically relevant VAT quantification.

3.5 VAT quantification

The VAT quantification module is designed to measure the volume of VAT accurately after segmentation and refinement using the U-Net. The goal is to calculate the VAT volume within the medical images, providing a quantitative assessment that can aid in diagnosing and monitoring obesity-related health risks.

Binary Mask Creation: The refined VAT regions from the previous steps are transformed into a binary mask, denoted as M . This mask assigns a value of 1 to pixels within the VAT region and 0 to pixels outside the VAT region. Mathematically, this is represented as:

$$M(i, j) = \begin{cases} 1 & \text{if pixel } (i, j) \text{ belongs to VAT region} \\ 0 & \text{otherwise} \end{cases} \quad (8)$$

where (i, j) are the coordinates of a pixel in the image. Figure 5 illustrates a simplified representation of the binary mask used in the VAT quantification algorithm.




Figure 5. Binary Mask Representation of VAT Segmentation

The blue region corresponds to the VAT mask (M), where each pixel is assigned a value of 1, indicating the presence of VAT. Conversely, the white region represents the non-VAT region, with each pixel assigned a value of 0.

Voxel Volume Calculation: Each pixel in the image corresponds to a three-dimensional voxel in the actual scanned volume. A voxel is a three-dimensional equivalent of a pixel. It has dimensions in width, height, and depth (usually denoted as $dx \times dy \times dz$). Each voxel in a medical image represents a small volume of tissue in the scanned body. The segmented VAT region consists of a set of connected voxels that have been identified by the U-Net and refined by the region-growing algorithm.

With an in-plane resolution of $r \times r$ (pixel width and height) and slice thickness t , the voxel volume is calculated as:

$$V_{\text{voxel}} = r^2 * t \quad (9)$$

VAT Volume Calculation: The total VAT volume, V_{VAT} , is calculated by summing the volumes of all voxels identified as VAT within the binary mask M . The algorithm first counts the number of pixels in the mask labelled as VAT (value 1), denoted as N_{VAT} . Then, the total VAT volume is calculated as:

$$V_{\text{VAT}} = N_{\text{VAT}} * V_{\text{voxel}} \quad (10)$$

Algorithm 2. VAT quantification

1. **Initialization:**
Set V_{VAT} to 0
Obtain the binary mask M from the refined segmentation.
 2. **Voxel Iteration:**
For each pixel (i, j) in M
If $M(i, j) == 1$
Increment N_{VAT}
 3. **Volume Calculation:**
Calculate V_{voxel} based on the image resolution (r, t)
Calculate $V_{\text{VAT}} = N_{\text{VAT}} * V_{\text{voxel}}$
 4. **Output:**
The final value of V_{VAT} represents the total volume of VAT in the image.
-

By integrating these steps, the VAT quantification module provides a comprehensive and reliable assessment of VAT volume, contributing valuable information for clinical diagnosis and treatment planning. Figure 6 explains the detailed process flow of the proposed VAT quantification system. Table 1 summarizes the list of symbols and notations in this article.

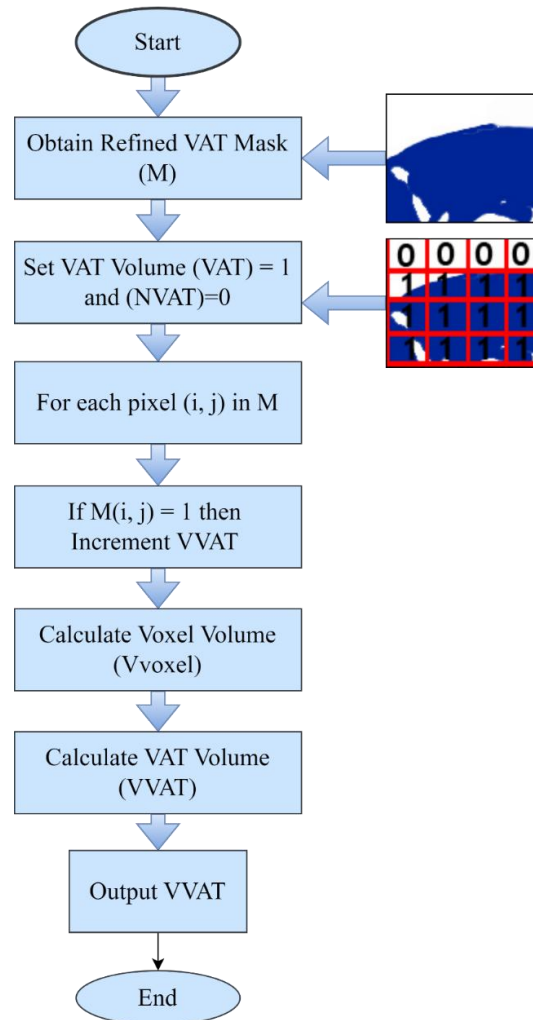


Figure 6. VAT quantification process flow.

Table 1. List of symbols and notations

Symbol/Notation	Description
$I(x, y)$	Observed image intensity at pixel location (x, y) .
$R(x, y)$	True image intensity at pixel location (x, y) .
$B(x, y)$	Bias field at pixel location (x, y) .
$F(u, v)$	Fourier transform of the logarithmic image.
$H(u, v)$	High-pass filter applied in the frequency domain.
$F_H(u, v)$	High-pass filtered Fourier transform.
$I_H(x, y)$	Corrected image after applying homomorphic filtering.
N_{VAT}	Number of voxels classified as VAT in the binary mask.
V_{voxel}	Volume of a single voxel based on image resolution and slice thickness.

V_{VAT}	Total VAT volume obtained from the segmented regions.
$M(i, j)$	Binary mask, where $M(i, j) = 1$ indicates a pixel belonging to VAT region.
r	Pixel width and height (in-plane resolution).
t	Slice thickness in the image volume.
ϵ	Predefined threshold for intensity similarity in the region-growing algorithm.
\log	Logarithmic transformation function.
\exp	Exponential function to revert the logarithmic transformation.
$\sigma(x)$	Sigmoid activation function used in U-Net for pixel classification.
F^{-1}	Inverse Fourier transform.

4. Results

The experimental setup for this research involved utilizing a Dell G15 computer equipped with an i5 processor, 16GB of RAM, and an NVIDIA QUADRO RTX 6000 graphics card with 24GB of VRAM. The operating system used was Windows 11.

During the training phase, 500 epochs were initially allocated, but it was observed that optimal performance was achieved within 300 epochs for all models. Following training, the segmentation efficiency of both the proposed and existing methods was assessed for VAT and non-VAT regions in abdominal MR images. Two distinct datasets were employed to thoroughly evaluate the performance of each method. A comparative analysis was conducted, including the proposed approach and various existing methods from the literature including Hao Shen et al. [15], Prakash et al [16], Kafali et al. [17], Somasundaram et al [18], Ogunleye et al. [19], Wu et al. [20], Christine Haugen et al. [21], Ina Vernikouskaya. [22], and Hwang et al. [23].

To ensure a comprehensive evaluation, multiple accuracy metrics were utilized, including the Dice Similarity Coefficient (DSC), Intersection over Union (IoU), Sensitivity, and Specificity. These metrics provide insights into the agreement between the predicted segmentation and the ground truth, as well as the ability to correctly identify VAT regions while minimizing false positives. The following equations are used to calculate the above performance metrics.

$$DSC = \frac{2 \cdot TP}{2 \cdot TP + FP + FN} \quad (11)$$

$$IoU = \frac{TP}{TP + FP + FN} \quad (12)$$

$$Sensitivity = \frac{TP}{TP + FN} \quad (13)$$

$$Specificity = \frac{TN}{TN + FP} \quad (14)$$

The hyperparameters for the proposed method were carefully selected. The training process used 500 epochs, although peak performance was reached before this limit. A learning rate of 0.0001 and a momentum value of 0.9 were employed to optimize the learning process. The Adam optimizer was chosen for its effectiveness in training deep learning models, particularly for image segmentation tasks.

4.1 k-Fold Cross-Validation

Tables 2 and 3 present a comprehensive assessment of the proposed method's performance using k-fold cross-validation on two distinct datasets: CHAOS and TCIA. Each dataset was evaluated across different folds (k=2, 3, 4, and 5), providing a robust evaluation of the method's consistency and generalization capabilities. For the CHAOS dataset, the method consistently achieved high DSC and IoU values across all folds, indicating excellent agreement between the predicted VAT segmentation and the ground truth. Sensitivity and specificity values also remained high, suggesting that the model effectively identified VAT regions while

minimizing false positives and false negatives. Notably, the performance improved as the number of folds increased, highlighting the method's ability to learn and generalize from the data. Similarly, on the TCIA dataset, the method demonstrated strong performance across all folds, with consistently high DSC and IoU values. Sensitivity and specificity remained high, further affirming the model's accuracy in identifying VAT regions and distinguishing them from other tissues.

Table 2: K-Fold Cross-Validation Performance of VAT Segmentation on CHAOS Dataset

k-Fold	DSC	IoU	Sensitivity	Specificity
K=2	0.923 ± 0.07	0.921 ± 0.19	0.954	0.959
K=3	0.951 ± 0.06	0.932 ± 0.12	0.961	0.973
K=4	0.961 ± 0.08	0.951 ± 0.12	0.985	0.977
K=5	0.987 ± 0.05	0.986 ± 0.19	0.981	0.993

Table 3: K-Fold Cross-Validation Performance of VAT Segmentation on TCIA Dataset

k-Fold	DSC	IoU	Sensitivity	Specificity
K=2	0.921 ± 0.09	0.933 ± 0.02	0.942	0.938
K=3	0.932 ± 0.11	0.941 ± 0.05	0.961	0.946
K=4	0.956 ± 0.21	0.957 ± 0.07	0.982	0.974
K=5	0.972 ± 0.16	0.963 ± 0.11	0.990	0.985

4.2 VAT segmentation efficiency analysis using different experimental scenarios

The experimental design aimed to thoroughly assess the impact of each component in the proposed method by evaluating four distinct scenarios. In the first scenario, both preprocessing and region growing were excluded, isolating the performance of the U-Net model alone. The second scenario excluded preprocessing but included region growing, highlighting the contribution of boundary refinement. The third scenario involved preprocessing without region growing, emphasizing the importance of image enhancement. Finally, the fourth scenario combined both preprocessing and region growing, representing the full proposed methodology. This approach allowed for a granular analysis of the contribution of each step to the overall performance. The results, presented as the median DSC for VAT across both datasets and 5-fold cross-validation, along with the global average, provided a comprehensive comparison.

Table 4 illustrates the impact of different components on VAT segmentation efficiency. The baseline U-Net model (Scenario 1) achieved reasonable performance, with median DSC values of 0.91 for CHAOS and 0.89 for TCIA datasets. However, the incorporation of region growing (Scenario 2) consistently improved the median DSC, highlighting its effectiveness in refining the segmentation results. Similarly, the addition of preprocessing (Scenario 3) led to notable improvements compared to the baseline, demonstrating the positive impact of image enhancement on segmentation accuracy.

Table 4: Comparative Analysis of VAT Segmentation Efficiency Across Different Experimental Scenarios and Datasets

Dataset	Scenario	Fold 1	Fold 2	Fold 3	Fold 4	Fold 5	Median DSC	Global Average
CHAOS	1	0.86	0.88	0.91	0.92	0.93	0.91	0.901
CHAOS	2	0.91	0.92	0.92	0.93	0.95	0.92	0.925
CHAOS	3	0.92	0.94	0.94	0.95	0.96	0.94	0.941
CHAOS	4	0.94	0.94	0.97	0.97	0.98	0.95	0.958
TCIA	1	0.85	0.88	0.89	0.90	0.91	0.89	0.886
TCIA	2	0.89	0.91	0.92	0.93	0.94	0.92	0.918
TCIA	3	0.90	0.92	0.94	0.95	0.96	0.94	0.922
TCIA	4	0.93	0.94	0.95	0.96	0.97	0.95	0.95

The full proposed method (Scenario 4), combining both preprocessing and region growing, consistently outperformed all other scenarios, showcasing the synergistic effect of these techniques in achieving the highest segmentation efficiency. These findings underscore the importance of each component in the proposed methodology, with the combination of preprocessing and region growth proving to be the most effective approach for accurate and robust VAT segmentation in abdominal MR images.

4.3 Model Performance Evaluation and Analysis

The charts (figure 7 a-d) illustrate the performance of a model during training and validation across various metrics. The training accuracy increases rapidly at the beginning and stabilizes quickly, indicating the model learns effectively from the training data. However, the validation accuracy, while initially rising, fluctuates significantly and doesn't reach the same level as the training accuracy, suggesting potential overfitting.

The IoU coefficient, a measure of overlap between the predicted and actual values, shows a similar trend. While the training IoU increases and then plateaus, the validation IoU exhibits noticeable fluctuations, indicating that the model might struggle to generalize to unseen data. The Dice coefficient, another metric assessing the similarity between predicted and actual values, mirrors the pattern observed in the IoU. This reinforces the notion that the model's performance on validation data isn't as robust as its performance on training data. Finally, the loss curves show that the training loss decreases steadily, as expected. However, the validation loss, while initially declining, later rises and fluctuates. This again points to the possibility of overfitting, where the model becomes too specialized to the training data and performs poorly on new, unseen data.

4.4 Model Performance Evaluation with Existing Methods

The state-of-the-art analysis for VAT segmentation models indicates varying levels of performance across different methods (Table 5). In order to get trustful results, every state-of-the-art U-net and CNN variant is trained and tested with the two mentioned datasets, and average values are taken for comparison. SECANet by Hao Shen et al. achieves a high sensitivity of 0.91 but struggles with IoU at 0.80, showing limitations in overlap accuracy. Similarly, the MultiRes-Attention U-Net from Prakash et al. performs well with a DSC of 0.87, though its IoU remains lower at 0.74. Other models like ACD 3D U-Net from Kafali et al. and nnU-Net by Somasundaram et al. face challenges with complex anatomical boundaries, resulting in sensitivities and specificities around 0.87-0.89. Ogunleye et al. propose a U-Net model with a relatively balanced performance, reaching an IoU of 0.85. Wu et al. with 2D-CDFNet and Christine Haugen et al.'s semi-automated approach struggle to exceed 0.88 in specificity. Notably, the proposed method outperforms existing models, with a DSC of 0.921, IoU of 0.955, and high sensitivity of 0.961, demonstrating its superior ability to address complex VAT segmentation challenges effectively.

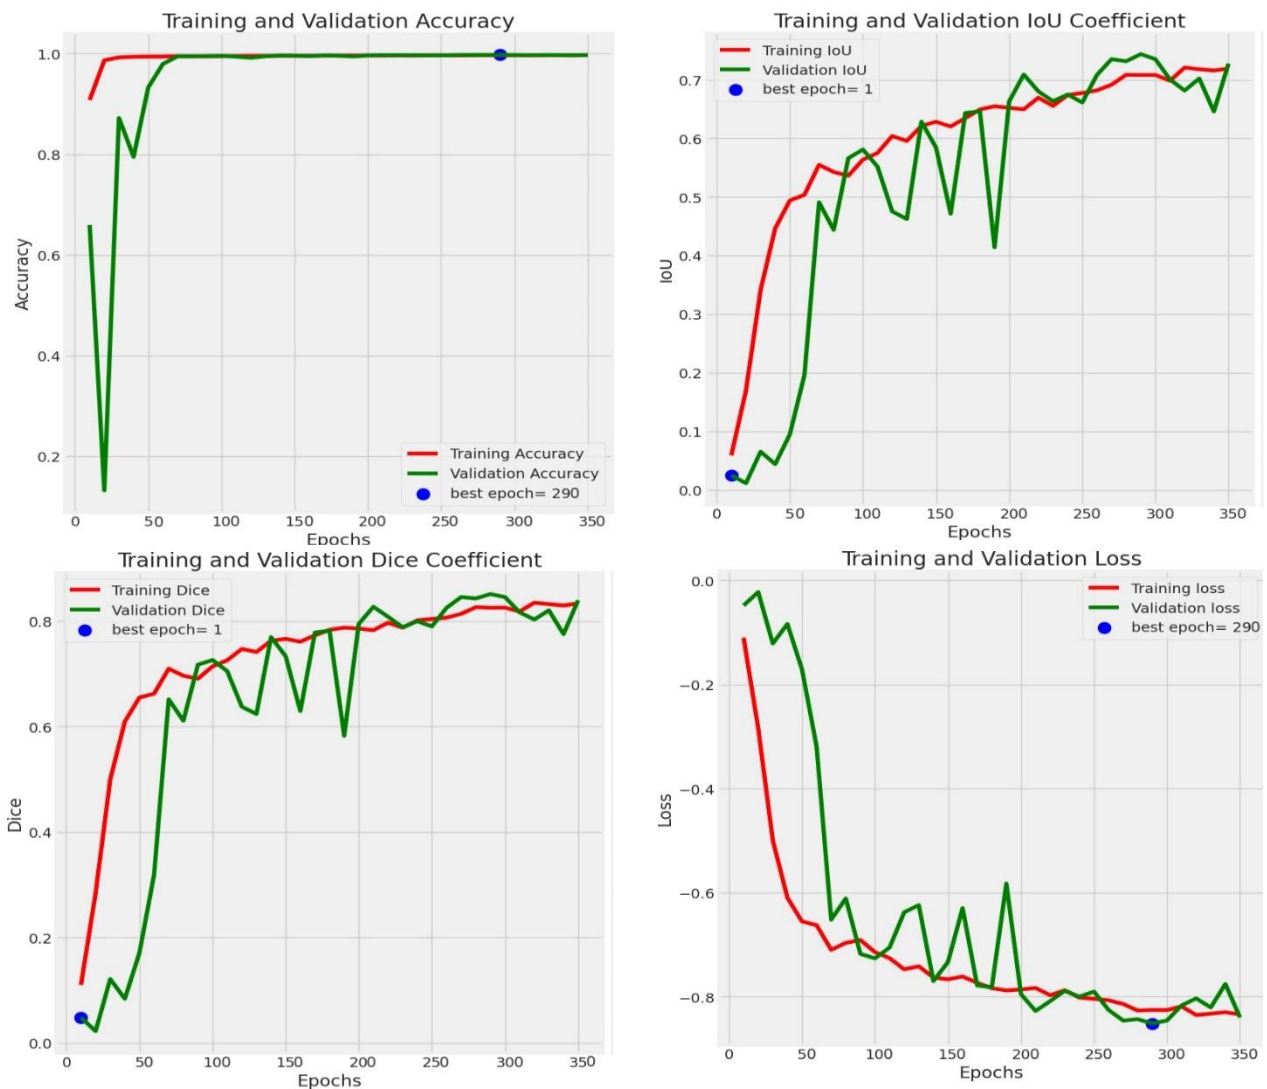


Figure 7 (a) Training and Validation Accuracy Over Epochs, **Figure 7 (b)** Training and Validation IoU Over Epochs, **Figure 7 (c)** Training and Validation Dice Coefficient Over Epochs, **Figure 7 (d)** Training and Validation Loss Over Epochs and Comparative analysis with existing methods.

Table 5: Comparison of VAT Segmentation Performance with Existing Methods

Author	Segmentation model name	DSC	IoU	Sensitivity	Specificity
Hao Shen et al. [15]	SECANet (based on U-Net)	0.87	0.80	0.91	0.88
Prakash et al [16]	MultiRes-Attention U-Net	0.87	0.74	0.89	0.81
Kafali et al. [17]	ACD 3D U-Net, 3D nnU-Net	0.85	0.73	0.87	0.79
Somasundaram et al [18]	nnU-Net	0.86	0.75	0.89	0.81
Ogunleye et al. [19]	U-Net	0.86	0.85	0.91	0.88
Wu et al. [20]	2D-CDFNet	0.84	0.75	0.86	0.81
Christine Haugen et al. [21]	Semi-Automated Approach	0.82	0.76	0.88	0.83
Ina Vernikouskaya et al. [22]	U-Net-like CNN	0.90	0.85	0.92	0.91
Hwang et al. [23]	Custom 3D CNN	0.86	0.81	0.91	0.87
Proposed method	Proposed	0.921	0.955	0.961	0.948

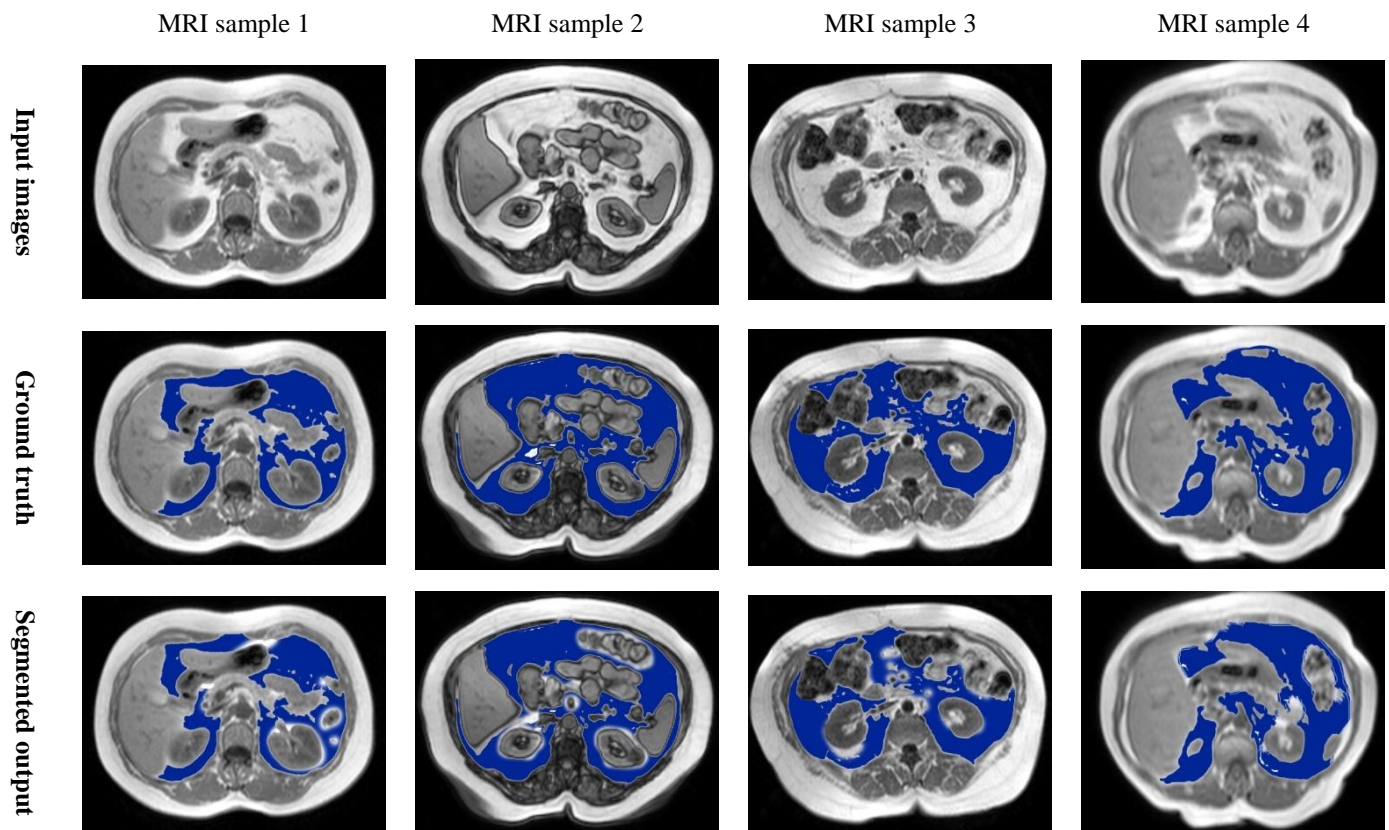


Figure 8: Visual Comparison of Ground Truth and Proposed Segmentation Method for VAT on Four Abdominal MR Image Samples

Figure 8 presents a visual comparison of the ground truth VAT segmentations (second row) with the corresponding segmentations produced by the proposed method (third row) for four abdominal MR image samples (first row). These samples are sourced from two different datasets, demonstrating the method's ability to generalize across varying image characteristics. The proposed method consistently demonstrates close alignment with the ground truth segmentations, accurately delineating the boundaries of VAT regions. This visual agreement is particularly evident in the third and fourth samples, where the method successfully captures the complex shapes and varying intensities of VAT. While minor discrepancies can be observed in some samples, such as a slight overestimation in the second sample, the overall results showcase the effectiveness of the proposed method in achieving precise VAT segmentation. This visual comparison not only confirms the quantitative results presented in previous sections but also highlights the qualitative strength of the proposed method in producing visually plausible and accurate segmentations of VAT in abdominal MR images.

4.5 Discussion

The experimental evaluation of VAT segmentation performance involved a comprehensive comparison of various existing methods against the proposed methodology. The comparison focused on key metrics like the DSC, IoU, sensitivity, and specificity, providing a multi-faceted assessment of segmentation accuracy.

The analysis revealed that the proposed method, utilizing an optimized U-Net architecture with region growing, consistently outperformed existing approaches across all metrics. This superior performance is attributed to several factors:

Bias Field Correction: The incorporation of homomorphic filtering for bias field correction effectively addressed intensity inhomogeneities in the MR images, enhancing the quality of input data and contributing to more accurate segmentation.

Optimized U-Net Architecture: The enhanced U-Net model, with its attention mechanisms and skip connections, proved adept at capturing subtle VAT patterns amidst complex anatomical structures. The attention mechanism allowed the model to focus on relevant features while suppressing noise, leading to improved boundary delineation.

Region Growing Refinement: The integration of region growing significantly enhanced the segmentation results by refining the boundaries of VAT regions identified by the U-Net. This step effectively addressed potential over- or under-segmentation issues, ensuring a more precise representation of the actual VAT distribution.

Diverse Datasets: The utilization of two distinct datasets, CHAOS and TCIA, contributed to the robustness of the proposed method. By training and evaluating diverse data, the model was able to generalize effectively to unseen cases, exhibiting high performance across a wider range of variations in VAT shape, size, and intensity.

The visual comparison in Figure 8 further strengthens the quantitative findings. The proposed method demonstrates close alignment with the ground truth segmentations, accurately delineating VAT boundaries in diverse MR image samples. While minor discrepancies are observed in some cases, the overall results highlight the method's ability to produce visually plausible and accurate segmentations.

However, the study acknowledges potential limitations. The performance of the proposed method may be affected by extreme variations in image quality or the presence of pathological conditions that alter the appearance of VAT. Future research could explore strategies to address these limitations and further enhance the robustness of the method.

5. Conclusion

The research addresses the critical problem of accurately segmenting and quantifying VAT from abdominal MRI images, which is essential for assessing obesity-related health risks. Traditional methods struggle with irregular shapes and intensity variations in VAT regions, leading to inaccurate results. The proposed methodology introduces a solution to these challenges through a three-step process: homomorphic filtering for intensity inhomogeneity correction, an optimized U-Net architecture with attention mechanisms for segmentation, and a region-growing algorithm for refining the segmentation. Two datasets were used in this study: the CHAOS dataset and TCIA dataset. These datasets provide diverse and challenging MRI scans, ensuring the model's ability to generalize and perform well on unseen data. The use of homomorphic filtering addresses the problem of intensity variations, which improves the quality of the input images. The enhanced U-Net architecture accurately segments VAT, even in cases with complex anatomical structures and varying intensities, while the region-growing algorithm refines the segmentation by focusing on VAT boundaries, addressing any over- or under-segmentation issues. The results of this work are highlighted by its superior performance in VAT segmentation. The method achieved high DSC values of 0.987 and 0.972 on the CHAOS and TCIA datasets, respectively. Similarly, the IoU values reached 0.986 and 0.963, indicating strong segmentation accuracy. Sensitivity values were 0.981 on the CHAOS dataset and 0.990 on the TCIA dataset, with specificity scores of 0.993 and 0.985, respectively. These results demonstrate a significant improvement over existing methods and provide a reliable tool for clinical applications in obesity-related health risk assessments.

Conflicts of Interest

The authors have no Conflicts of Interest

References

- [1] M. Shuster, M. Patlas, J. H. Pinthus, and M. Mourtzakis, "The Clinical Importance of Visceral Adiposity: A Critical Review of Methods for Visceral Adipose Tissue Analysis," *The British Journal of Radiology*, Vol.85, No.1009, pp.1-10, Jan. 2012.
- [2] A. Chait and L. J. den Hartigh, "Adipose Tissue Distribution, Inflammation and Its Metabolic Consequences, Including Diabetes and Cardiovascular Disease," *Frontiers in Cardiovascular Medicine*, Vol.7, No.22, Feb. 2020.

- [3] C. Storz, S. D. Heber, S. Rospleszcz, J. Machann, S. Sellner, K. Nikolaou, R. Lorbeer, S. Gatidis, S. Elser, A. Peters, C. L. Schlett, and F. Bamberg, "The Role of Visceral and Subcutaneous Adipose Tissue Measurements and Their Ratio by Magnetic Resonance Imaging in Subjects with Prediabetes, Diabetes and Healthy Controls from a General Population without Cardiovascular Disease," *Br J Radiol*, Vol.91, No.1089, pp.20170808, 2018.
- [4] M. Małodobra-Mazur, A. Cierznia, Dorota Pawełka, K. Kaliszewski, J. Rudnicki, and T. Dobosz, "Metabolic Differences between Subcutaneous and Visceral Adipocytes Differentiated with an Excess of Saturated and Monounsaturated Fatty Acids," *Genes*, vol. 11, no. 9, pp. 1092–1092, Sep. 2020
- [5] H. Liu, D. Yang, S. Li, Y. Xiao, Y. Tu, D. Peng, Y. Bao, J. Han, and H. Yu, "A Reliable Estimate of Visceral Fat Area from Simple Anthropometric Measurements in Chinese Overweight and Obese Individuals," *Front Endocrinol (Lausanne)*, Vol.13, pp.916124, Jun. 2022
- [6] L. C. Pescatori, E. Savarino, G. Mauri, E. Silvestri, M. Cariati, F. Sardanelli, and L. M. Sconfienza, "Quantification of Visceral Adipose Tissue by Computed Tomography and Magnetic Resonance Imaging: Reproducibility and Accuracy," *Radiol Bras*, Vol.52, No.1, pp.1-6, Jan-Feb. 2019.
- [7] H. H. Hu, J. Chen, and W. Shen, "Segmentation and quantification of adipose tissue by magnetic resonance imaging," *Magnetic Resonance Materials in Physics, Biology and Medicine*, vol. 29, no. 2, pp. 259–276, Sep. 2015
- [8] Shen N, Li X, Zheng S, Zhang L, Fu Y, Liu X, Li M, Li J, Guo S, Zhang H. Automated and accurate quantification of subcutaneous and visceral adipose tissue from magnetic resonance imaging based on machine learning. *Magn Reson Imaging*. 2019 Dec;64:28-36. doi: 10.1016/j.mri.2019.04.007. Epub 2019 Apr 18. PMID: 31004712.
- [9] B. T. Addeman, S. Kutty, T. G. Perkins, A. S. Soliman, C. N. Wiens, C. M. McCurdy, M. D. Beaton, R. A. Hegele, and C. A. McKenzie, "Validation of Volumetric and Single-Slice MRI Adipose Analysis Using a Novel Fully Automated Segmentation Method," *J. Magn. Reson. Imaging*, Vol.41, pp.233-241, 2015.
- [10] A. A. Joshi, H. H. Hu, R. M. Leahy, M. I. Goran, and K. S. Nayak, "Automatic Intra-Subject Registration-Based Segmentation of Abdominal Fat from Water–Fat MRI," *J. Magn. Reson. Imaging*, Vol.37, 2013.
- [11] D. Wald, B. Teucher, J. Dinkel, R. Kaaks, S. Delorme, H. Boeing, K. Seidensaal, H.-P. Meinzer, and T. Heimann, "Automatic Quantification of Subcutaneous and Visceral Adipose Tissue from Whole-Body Magnetic Resonance Images Suitable for Large Cohort Studies," *J. Magn. Reson. Imaging*, Vol.36, pp.1421-1434, 2012.
- [12] C. Haugen, V. Lysne, I. Haldorsen, et al., "MRI Adipose Tissue Segmentation and Quantification in R (RAdipoSeg)," *Diabetol Metab Syndr*, Vol.14, pp.146, 2022.
- [13] G. Varoquaux and V. Cheplygina, "Machine Learning for Medical Imaging: Methodological Failures and Recommendations for the Future," *NPJ Digit Med*, Vol.5, No.1, pp.48, Apr. 2022.
- [14] G.-S. Fu, Y. Levin-Schwartz, Q.-H. Lin, and D. Zhang, "Machine Learning for Medical Imaging," *Journal of Healthcare Engineering*, Vol.2019, Article ID 9874591, 2 pages, 2019.
- [15] H. Shen, Y. Ren, Z. Huang, G. Li S Wang, M. Cong, D. Luo, D. Shao, EY. Lee, L. Cui R Huo, J. Qin, J. Liu, Z. Hu, Z. Liu, and N. Zhang, "A deep learning model based on the attention mechanism for automatic segmentation of abdominal muscle and fat for body composition assessment," *Quantitative imaging in medicine and surgery*, vol. 13, no. 3, pp. 1384–1398, Mar. 2023
- [16] B. Prakash, A. Channarayana Srinivasa, L. Yun Yeow, W. Xiang Chen, A. Jing, W. Shiong Lim, and C. Heng Tan, "MultiRes Attention Deep Learning Approach for Abdominal Fat Compartment Segmentation and Quantification," *Artificial intelligence*, Nov. 2023.
- [17] S. Kafali, Shih, X. Li, G. Hyun, T. Kelly, S Chowdhury, S Loong, J Moretz, S.R. Barnes, Z. Li, and H. H. Wu, "Automated abdominal adipose tissue segmentation and volume quantification on longitudinal MRI using 3D convolutional neural networks with multi-contrast inputs," *Magnetic Resonance Materials in Physics Biology and Medicine*, Feb. 2024
- [18] A. Somasundaram, M. Wu, A. Reik, S. Rupp, J. Han, S. Naebauer, D. Junker, L. Patzelt, M. Wiechert, Y. Zhao, D. Rueckert, H. Hauner, C. Holzappel, and D. C. Karampinos, "Evaluating Sex-specific Differences in Abdominal Fat Volume and Proton Density Fat Fraction on MRI Scans Using Automated nnU-Net-based Segmentation," *Radiology Artificial Intelligence*, Vol.6, No.4, 2024.
- [19] O. A. Ogunleye, H. Raviprakash, A. M. Simmons, R. T. M. Bovell, P. E. Martinez, J. A. Yanovski, K. F. Berman, P. J. Schmidt, E. C. Jones, H. Bagheri, N. M. Biassou, and L. Y. Hsu, "A Combined Region- and Pixel-Based Deep Learning Approach for Quantifying Abdominal Adipose Tissue in Adolescents Using Dixon Magnetic Resonance Imaging," *Tomography*, Vol.9, No.1, pp.139-149, 2023.
- [20] T. Wu, S. Estrada, R. van Gils, R. Su, V. W. V. Jaddoe, E. H. G. Oei, and S. Klein, "Automated Deep Learning-Based Segmentation of Abdominal Adipose Tissue on Dixon MRI in Adolescents: A Prospective Population-Based Study," *AJR Am J Roentgenol*, Vol.222, No.1, pp.e2329570, 2024.

- [21] C. Haugen, V. Lysne, I. Haldorsen, E. Tjora, O. A. Gudbrandsen, J. V. Sagen, S. N. Dankel, and G. Mellgren, "MRI Adipose Tissue Segmentation and Quantification in R (RADipoSeg)", *Diabetol Metab Syndr*, Vol.14, No.1, pp.146, 2022
- [22] I. Vernikouskaya, H. P. Müller, D. Felbel, F. Roselli, A. C. Ludolph, J. Kassubek, and V. Rasche, "Body fat compartment determination by encoder-decoder convolutional neural network: application to amyotrophic lateral sclerosis," *Scientific Reports*, vol. 12, no. 1, p. 5513, Apr. 2022
- [23] J. J. Hwang and K. Pak, "Development of Automated Segmentation of Visceral Adipose Tissue in Computed Tomography," *European Journal of Radiology*, Vol.157, pp.110559–110559, Dec. 2022.
- [24] Z. Wang, Y. Meng, F. Weng, Y. Chen, F. Lu, X. Liu, M. Hou, and J. Zhang, "An Effective CNN Method for Fully Automated Segmenting Subcutaneous and Visceral Adipose Tissue on CT Scans," *Ann Biomed Eng*, Vol.48, No.1, pp.312-328, Jan. 2020.
- [25] K. Kaur, N. Jindal, and K. Singh, "Improved Homomorphic Filtering Using Fractional Derivatives for Enhancement of Low Contrast and Non-Uniformly Illuminated Images," *Multimedia Tools and Applications*, Vol.78, No.19, pp.27891–27914, Jun. 2019.
- [26] P. Sharma, T. Ninomiya, K. Omodaka, N. Takahashi, T. Miya, N. Himori, T. Okatani, and T. Nakazawa, "A Lightweight Deep Learning Model for Automatic Segmentation and Analysis of Ophthalmic Images," *Sci Rep*, Vol.12, No.1, pp.8508, May 2022.
- [27] CHAOS: Combined Healthy Abdominal Organ Segmentation Challenge. Available at: <https://chaos.grand-challenge.org/Data/>.
- [28] The Cancer Imaging Archive (TCIA): Clark, K., Vendt, B., Smith, K., Freymann, J., Kirby, J., Koppel, P. and Prior, F. (2013). The Cancer Imaging Archive (TCIA): Maintaining and Operating a Public Information Repository. *Journal of Digital Imaging*, 26(6), 1045-1057. Available at: <https://www.cancerimagingarchive.net/>