Lower Limb Calf Muscle Segmentation from Diffusion-Weighted Magnetic Resonance Images Using Deep Learning

Eshan Pandey [a], Xiaomeng Wang [b], Julian Gan [c], Ying-Hwey Nai [d], Derek J. Hausenloy [b,e,f,g], Pek Lan Khong [d], Forest Su Lim Tan [h], Thiruneepan Selvakulasingam [h], Ryan Fraser Kirwan [h] and Cheryl Pei Ling Lian [a]

[a] Health and Social Sciences Cluster, Singapore Institute of Technology, Singapore
E-mail: pandey.eshan@singaporetech.edu.sg
cheryl.lian@singaporetech.edu.sg

[b] Cardiovascular and Metabolic Disorders Program, Duke-National University of Singapore, Medical School, Singapore
E-mail: xiaomeng.wang@u.duke.nus.edu

[c] Siemens Healthineers, Singapore
E-mail: julian.gan@siemens-healthineers.com

[d] Clinical Imaging Research Centre, Yong Loo Lin School of Medicine, National University of Singapore, Singapore
E-mail: yinghweynai@yahoo.com
pek-lan.khong@nus.edu.sg

[e] National Heart Research Institute Singapore, National Heart Centre, Singapore

[f] Yong Loo Lin, School of Medicine, National University Singapore, Singapore

[g] The Hatter Cardiovascular Institute, University College London, London, UK
d.hausenloy@ucl.ac.uk

[h] Infocomm Technology Cluster, Singapore Institute of Technology, Singapore
forest.tan@singaporetech.edu.sg
thiru@singaporetech.edu.sg
ryan.kirwan@singaporetech.edu.sg

Abstract— **Peripheral artery disease (PAD) affects blood flow to the limbs, and diffusion-weighted magnetic resonance imaging (DW-MRI) can quantify microvascular perfusion and diffusion in calf muscles, aiding diagnosis. However, manual segmentation is time consuming and subjective. We propose a conditional generative adversarial network (cGAN) with an enhanced U-Net architecture for automated segmentation of calf muscles from DW-MRI. Our method leverages data augmentation to address small dataset sizes, splitting images into left and right halves and applying flipping and progressive rotation. Evaluated on datasets of healthy and PAD patients, our approach achieves average Dice Similarity Coefficient (Dice) scores of 54.86% to 79.85% across muscle groups, significantly outperforming baseline models (original U-Net architecture with cGAN and no data augmentation). This work demonstrates the potential of deep learning for automating segmentation in PAD diagnosis, offering a scalable solution for clinical applications.**

I. Introduction

Peripheral artery disease (PAD) is a prevalent condition characterized by narrowed arteries, reducing blood flow to the limbs and leading to inflammation-induced endothelial dysfunction impacting more than 200 million adults globally, with its occurrence rising to approximately 20% in individuals older than 70 years [1]. Conventional diagnostic methods, such as ankle-brachial index (ABI) and toe-brachial index (TBI), primarily assess macrovascular obstruction but fail to capture microvascular changes critical for understanding disease progression. Diffusion-weighted magnetic resonance imaging (DW-MRI) enables quantification of diffusion and perfusion in calf muscles, providing valuable insights into vascular health. However, manual segmentation of muscle regions from anatomical T1-weighted images is labour-intensive and prone to variability, necessitating automated solutions.

Funding: This study was supported by the National Research Foundation Competitive Research Program [NRF CRP25-2020RS-0001] and funded by the Singapore Institute of Technology (SIT) Academic Research Fund Tier 1 grant [WBS R-R12-A404-0002]. DJH is supported by the CArdiovascular DiseasE National Collaborative Enterprise (CADENCE) National Clinical Translational Program (MOH-001277-01).

979-8-3315-7206-8/25/$31.00 ©2025 IEEE

Figure 1. Diffusion-weighted MR image (left), manually delineated ROIs (centre), manually delineated ROIs overlaid on the DWI (right) of one subject

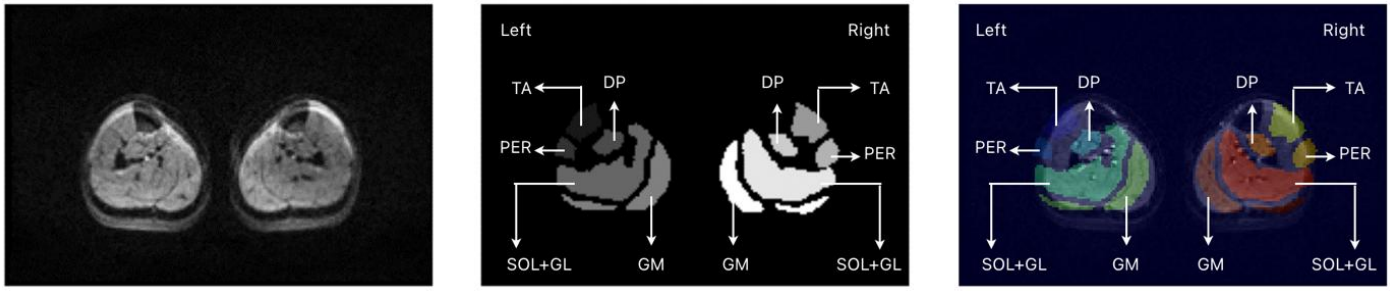

Deep learning, particularly convolutional neural networks (CNNs), has shown promise in medical image segmentation. In this work, we propose a conditional generative adversarial network (GAN) with an enhanced U-Net architecture to automate the segmentation of five calf muscle groups per limb: Tibialis anterior (TA), Peroneal muscles (PER), deep posterior compartment (DP), Gastrocnemius medialis (GM), and Soleus combined with Gastrocnemius lateralis (SOL+GL). The segmentation pipeline was implemented using our publicly available code [2]. Our contributions include:

• An enhanced U-Net architecture preserving spatial information and reducing original U-Net parameters.

• A data augmentation strategy involving image splitting, flipping, and rotation to improve model generalization.

• Comprehensive evaluation on healthy and PAD patient test datasets, achieving high Dice scores.

Figure 1 shows a sample DW-MRI and manually delineated calf muscle groups serving as the ground truth segmentations.

II. Related Work

Several studies have demonstrated the effectiveness of deep learning for muscle segmentation in MRI. For example, research has shown that convolutional neural networks (CNNs) can accurately segment cervical spine muscles from Dixon MRI [3], achieving high reliability. Similarly, U-Net models have been used to segment thigh muscles from fat-water decomposition MRI, with Dice scores above 0.85, indicating robust performance [4]. For lower leg muscles, de Bruin et al. (2021) [5] applied deep learning to segment 11 lower leg muscles and two bones from T1-weighted MRI scans of children with and without cerebral palsy. Their hybrid model, H-DenseUNet, achieved a Dice similarity coefficient of 0.90, showcasing the effectiveness of deep learning for complex segmentation tasks in the lower leg. Qureshi et al (2023) [6] explored deep learning for segmenting extraocular muscles from MRI scans, comparing four frameworks (U-Net, U-NeXt, DeepLabV3+, and ConResNet). The U-Net model [7] achieved the highest Dice score (0.85) underscoring the robustness of U-Net-based architectures for muscle segmentation. These studies highlight the potential of deep learning for muscle segmentation.

Despite the success of deep learning in muscle segmentation, its application to DW-MRI for calf muscle segmentation in PAD patients remains underexplored. Existing studies primarily utilize T1-weighted, Dixon, or Contrast-enhanced MRI (CE-MRI), which differ from DW-MRI's ability to quantify diffusion and perfusion separately. Moreover, while deep learning has been applied to vascular segmentation [8] and perfusion analysis in PAD [9], there is a lack of automated methods for segmenting calf muscles, which are critical for assessing microvascular health.

Khagi et al (2025) [10] investigated contrast-enhanced MRI (CE-MRI) to assess calf muscle perfusion in PAD patients, and reported that hypo-enhanced voxel features, indicative of impaired perfusion, across five muscle compartments (anterior, lateral, deep posterior, soleus, and gastrocnemius) effectively distinguished PAD patients from healthy controls using a decision tree classifier. While Khagi et al's study did not employ deep learning, it establishes the clinical importance of muscle perfusion analysis in PAD, which our work extends by automating segmentation using a conditional GAN [11] with enhanced U-Net architecture.

III. Methods

A. Data Acquisition

We employed two distinct datasets comprising lower-limb calf muscle diffusion-weighted MRI (DW-MRI) scans. Dataset 1 included scans from 24 healthy individuals. Each participant underwent both baseline and high-quality DW-MRI acquisitions, with test-retest imaging sessions within the same day. Baseline scans included 16 b-values, while high-quality scans contained either 10 or 16 b-values. Dataset 2 comprised scans from 9 patients diagnosed with peripheral artery disease (PAD), each of whom underwent a single baseline high-quality scan with 10 b-values. Across both datasets, b-values ranged from 0 to 1200 s/mm².

B. Data Annotation

Five regions of interest (ROIs) per calf (TA, PER, DP, GM, SOL GL) were manually delineated on structural MR images by undergraduate radiography students who were overseen by a senior medical student with 3 years of experience in image segmentation on MRI, avoiding tibia, fibula, and major blood vessels. This resulted in 10 ROIs per image, with class labels

1-5 for the left calf and 6-10 for the right calf after preprocessing.

C. Data Processing and Augmentation

To increase training data, DW-MRI images were split into left and right halves, doubling the sample size and simplifying segmentation to five classes per half. Each half was resized to 256x256 with a single channel. Data augmentation included flipping along x and y axes and rotating each image by 15-degree increments to generate 24 additional samples per image. ROIs were hot-encoded into six-channel masks (one for each muscle group and one for background).

D. Model Architecture

We employed a conditional GAN with an enhanced U-Net as the generator and a discriminator to distinguish predicted from manual masks. The enhanced U-Net adds forward blocks with skip connections after each upsampling and downsampling layer, preserving spatial information by not downsampling below 8x8 (unlike the original U-Net's 1x1). This reduces parameters (159.38 MB vs. 207.62 MB) while improving generalization. Forward block is an additional convolutional block structurally identical to the preceding convolution layer.

Table 1: Comparison of Original U-Net and proposed Enhanced U-Net architecture used for segmentation.

Original U-Net			Enhanced U-Net		
Layer	Output Shape	Params	Layer	Output Shape	Params
Input Layer	(None, 256, 256, 3)	0	Input Layer	(None, 256, 256, 1)	0
Sequential_2	(None, 128, 128, 64)	3,072	Sequential_2	(None, 128, 128, 64)	1,280
Sequential_3	(None, 64, 64, 128)	131,584	Sequential_3	(None, 128, 128, 64)	65,792
Sequential_4	(None, 32, 32, 256)	525,312	Sequential_4	(None, 64, 64, 128)	131,584
Sequential_5	(None, 16, 16, 512)	2,099,200	Sequential_5	(None, 64, 64, 128)	262,656
Sequential_6	(None, 8, 8, 512)	4,196,352	Sequential_6	(None, 32, 32, 256)	525,312
Sequential_7	(None, 4, 4, 512)	4,196,352	Sequential_7	(None, 32, 32, 256)	1,049,600
Sequential_8	(None, 2, 2, 512)	4,196,352	Sequential_8	(None, 16, 16, 512)	2,099,200
Sequential_9	(None, 1, 1, 512)	4,196,352	Sequential_9	(None, 16, 16, 512)	4,196,352
Sequential_10	(None, 2, 2, 512)	4,196,352	Sequential_19	(None, 8, 8, 1024)	8,392,704
Concatenate	(None, 2, 2, 1024)	0	Sequential_11	(None, 16, 16, 512)	8,390,656
Sequential_11	(None, 4, 4, 512)	8,390,656	Concatenate	(None, 16, 16, 1024)	0
Concatenate_1	(None, 4, 4, 1024)	0	Sequential_12	(None, 16, 16, 512)	8,390,656
Sequential_12	(None, 8, 8, 512)	8,390,656	Concatenate_1	(None, 16, 16, 1024)	0
Concatenate_2	(None, 8, 8, 1024)	0	Sequential_13	(None, 32, 32, 256)	4,195,328
Sequential_13	(None, 16, 16, 512)	8,390,656	Concatenate_2	(None, 32, 32, 512)	0
Concatenate_3	(None, 16, 16, 1024)	0	Sequential_14	(None, 32, 32, 256)	2,098,176
Sequential_14	(None, 32, 32, 256)	4,195,328	Concatenate_3	(None, 32, 32, 512)	0
Concatenate_4	(None, 32, 32, 512)	0	Sequential_15	(None, 64, 64, 128)	1,049,088
Sequential_15	(None, 64, 64, 128)	1,049,088	Concatenate_4	(None, 64, 64, 256)	0
Concatenate_5	(None, 64, 64, 256)	0	Sequential_16	(None, 64, 64, 128)	524,800
Sequential_16	(None, 128, 128, 64)	262,400	Concatenate_5	(None, 64, 64, 256)	0
Concatenate_6	(None, 128, 128, 128)	0	Sequential_17	(None, 128, 128, 64)	262,400
Conv2DTranspose_8	(None, 256, 256, 3)	6,147	Concatenate_6	(None, 128, 128, 128)	0
			Sequential_18	(None, 128, 128, 64)	131,328
			Concatenate_7	(None, 128, 128, 128)	0
			Conv2DTranspose_8	(None, 256, 256, 6)	12,294
Total params: 54,425,859 (207.62 MB) Trainable params: 54,414,979 (207.58 MB) Non-trainable params: 10,880 (42.50 KB)			Total params: 41,779,206 (159.38 MB) Trainable params: 41,769,478 (159.34 MB) Non-trainable params: 9,728 (38.00 KB)		

Table 1 shows a comparison of the original U-Net and the proposed Enhanced U-Net. To the best of our knowledge at the current time, there are no reported studies that have segmented lower limb calf muscles based on pre-trained DW-MR images. We started out by pretraining our baseline models. Our baseline models were cGAN models with original U-Net architecture and no data augmentation. SegMode_01 and SegMoel_02 were our baseline models. With proper data augmentation and model architecture optimization we observed progressive improvement in Dice score from SegMode_01 and SegMoel_02 to SegModel_05, SegModel_06 and SegModel_07. Table 3 summarizes the average Dice score calculated across all available b-value images for the test set for all the devised models.

E. Training Procedure

The model was trained to minimize binary cross-entropy (BCE) and mean absolute error (MAE) between predicted and ground-truth masks, with an additional adversarial loss for the GAN. Training used a batch size of 32, Adam optimizer with a learning rate 0.001, and 100 epochs. Early Stopping was used during the training. Data augmentation prevented overfitting by increasing sample diversity.

For SegModel_01 to SegModel_05 the objective of the network was to minimize the binary cross entropy (L_{BCE}) and mean average error (L_{MAE}) between the predicted and ground truth masks. The generator loss and discriminator loss were defined as follows:

$$L_{BCE}(Y,Z) = = \frac{1}{B}\sum_{i=1}^{B} -\left\{Y^i * \log\left(\sigma(Z^i)\right) + (1 - Y^i) * \log\left(1 - \sigma(Z^i)\right)\right\} \quad (1)$$

$$L_{MAE}(Y,Z) = \left(\frac{1}{B}\sum_{i=1}^{B}\left|Y^i - Z^i\right|\right) \quad (2)$$

$$Generator\ Loss = L_{BCE}(Y_{True}, Y_{Pred}) + \lambda * L_{MAE}(Y_{True}, Y_{Pred}) \quad (3)$$

$$Discriminator\ Loss = L_{BCE}(Y_{True}, Tens(1)) + L_{BCE}(Y_{Pred,} Tens(0)) \quad (4)$$

$\sigma(\,) = Sigmoid\ function;\ \lambda = 100;$

$Tens(a)$ *is tensor with all elements* a.

SegModel_01 and SegModel_02 (baseline models) used a cGAN framework with the original U-Net architecture. SegModel_01 was trained solely on healthy subjects, while SegModel_02 included both healthy and PAD patient data. SegModel_03 and SegModel_04 employed an enhanced U-Net architecture, with SegModel_03 trained on healthy subjects only, and SegModel_04 on the combined dataset. None of these models used data augmentation. SegModel_05, which applied augmentation as described in Section III.C, served as the base model for SegModel_06 and SegModel_07. SegModel_06 was fine-tuned on healthy subjects, and SegModel_07 on both healthy and PAD subjects. During fine-tuning, only the enhanced U-Net generator was updated (the discriminator was removed), and training minimized BCE and channel-wise Dice loss. No data augmentation was applied during fine-tuning; it was limited to the initial training phase.

$$Dice\ coef(Y,Z) = 2\sum_{i=1}^{B} Y^i * \sigma(Z^i) \Big/ \sum_{i=1}^{B} Y^i + \sum_{i=1}^{B} \sigma(Z^i) \quad (5)$$

$$L_{Dice}(Y,Z) = 1 - Dice\ coef(Y,Z) \quad (6)$$

Combined loss for finetuning then is:

$$L_{combined} = \propto L_{BCE} + (1{-}\propto) * L_{Dice} \quad (7)$$

$\propto\ is\ 0.5$

IV. Experiments And Results

A. Experimental Setup

Dataset 1 (healthy subjects) was split into 18 for training, and 3 each for validation and testing. Dataset 2 (PAD patients) contributed 5 for training, and 2 each for validation and testing.

Each healthy subject underwent four scans: baseline test, baseline retest, high-quality test, and high-quality retest. Set 1 included 16 b-value scans for baseline test/retest and 10 for high-quality test/retest. Set 2 included 16 b-value scans for all four scan types. PAD patients had only baseline and high-quality tests, with 16 and 10 b-value scans, respectively. All b-value images were used in training and evaluation.

Figure 2 Dataset split. All b-values are expressed in s/mm2

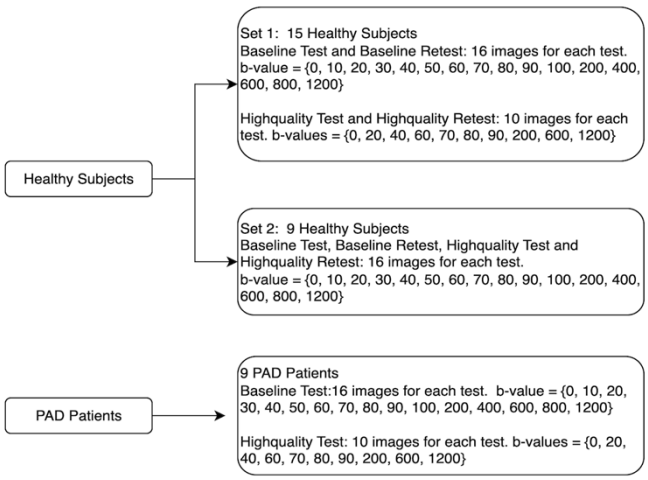

Some models were trained solely on healthy data, while others used combined healthy and PAD datasets. Regardless of training composition, all models were evaluated on the full test set across all b-values. Table 2 outlines the configurations.

Table 2 Training configuration of various models

MODELS	Discriminator	Generator	Training Data	Data Augmentation	Fine tuning
SegModel_0 (Baseline)	Yes	Original U-Net	Healthy Subjects	No	No
SegModel_02 (Baseline)	Yes	Original U-Net	All Subjects	No	No
SegModel_03	Yes	Enhanced U-Net	Healthy Subjects	No	No
SegModel_04	Yes	Enhanced U-Net	All Subjects	No	No
SegModel_05	Yes	Enhanced U-Net	Healthy Subjects	Yes	No
SegModel_06	No	Enhanced U-Net	Healthy Subjects	Yes	Yes
SegModel_07	No	Enhanced U-Net	All Subjects	Yes	Yes

B. Results

Table 3 presents the average Dice scores per muscle group across models. SegModel_07 achieved the highest scores (54.86%–79.85%), demonstrating the impact of data

augmentation and fine-tuning. Example segmentations in Fig. 3 highlight improved boundary delineation with SegModel_07.

We employ a conditional Generative Adversarial Network (cGAN) with a U-Net generator to segment two image sets: healthy subjects and patients with Peripheral Artery Disease (PAD). To address the limitations of standard U-Net, we introduce an Enhanced U-Net architecture. Seven models (SegModel_01–07) were evaluated using various combinations of network design, dataset composition, augmentation, and fine-tuning protocols (Table 2).

Baseline performance was established using SegModel_01 (U-Net, healthy-only) and SegModel_02 (U-Net, healthy + PAD). Both performed suboptimally, with SegModel_01 unexpectedly achieving a higher Dice score despite a smaller, more homogeneous training set. We hypothesise that limited anatomical variation among healthy subjects enabled the model to converge to a simplistic local minimum. In contrast, the greater structural variability within the PAD cohort posed a more complex optimisation problem.

SegModel_03 (Enhanced U-Net, healthy-only) underperformed; however, SegModel_04 (Enhanced U-Net, healthy + PAD) yielded more consistent scores across muscle groups, in contrast to the high variance seen in SegModel_01. This suggests that the enhanced architecture, when exposed to diverse anatomy, converges to a more robust and generalisable local minimum.

SegModel_05–07 represent our key contributions, leveraging the enhanced architecture with refined training strategies to significantly improve segmentation accuracy. While Dice scores remain modest, it is important to note that the dataset comprised only 33 subjects (23 for training). Despite this limitation, results are promising and indicate that larger training sets would likely yield further gains in segmentation performance.

Figure 3 Visual comparison of the manually drawn region of interest and auto segmentation (predicted region of interest) using model SegModel_07

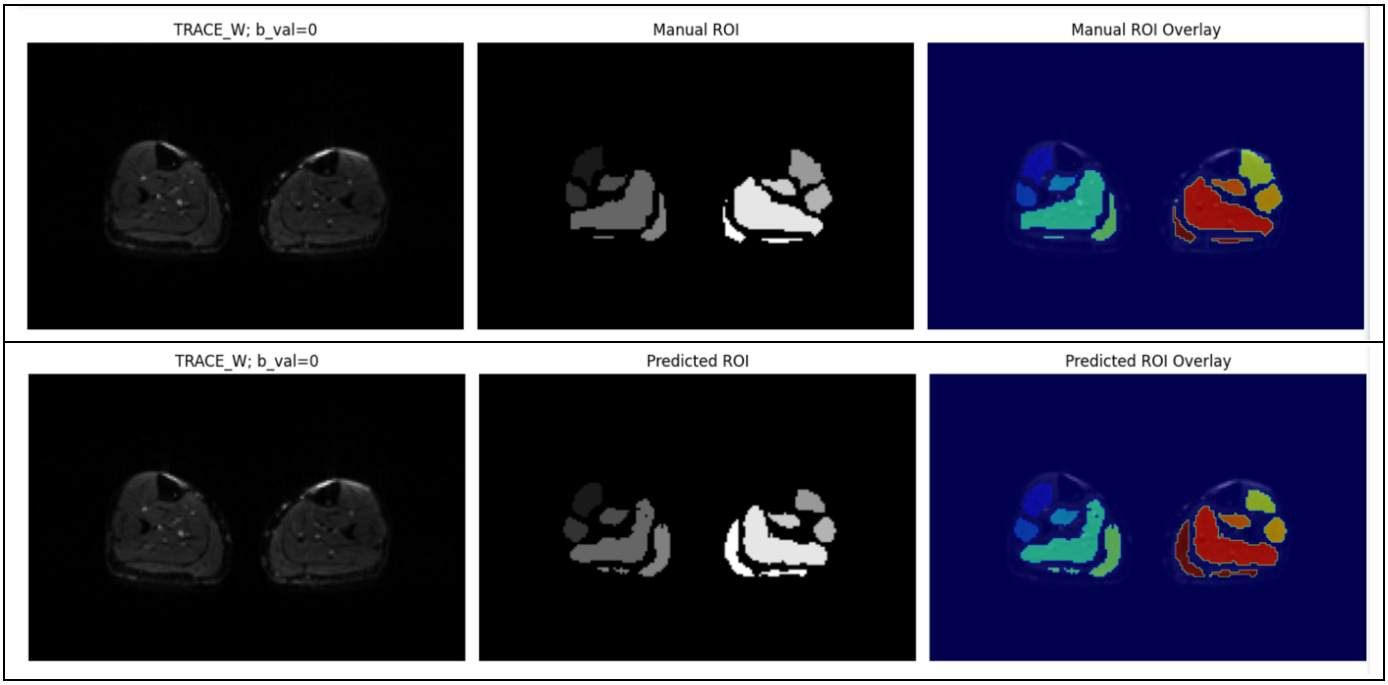

Table 3 Average Dice score for each muscle group by various segmentation models

MODELS	TA left	PER left	DP left	SOL GL Left	GM left	TA right	PER right	DP right	SOL GL right	GM right
SegModel_01 (baseline)	29.32%	4.09%	1.01%	4.26%	2.65%	29.23%	4.22%	1.61%	5.11%	1.88%
SegModel_02 (baseline)	2.27%	1.21%	2.86%	0.00%	0.00%	10.88%	2.07%	1.74%	0.00%	0.00%
SegModel_03	1.98%	0.14%	0.21%	1.24%	0.19%	0.89%	0.08%	1.95%	1.95%	0.22%
SegModel_04	9.03%	2.40%	4.20%	0.00%	7.21%	7.68%	3.21%	3.83%	0.16%	7.65%
SegModel_05	62.42%	34.81%	36.12%	65.93%	47.47%	67.86%	44.98%	42.99%	70.95%	48.42%
SegModel_06	38.54%	27.48%	51.71%	67.37%	45.84%	69.57%	56.52%	51.65%	71.61%	50.18%
SegModel_07	**66.89%**	**54.86%**	**66.24%**	**77.82%**	**57.95%**	**79.85%**	**74.27%**	**62.83%**	**78.04%**	**60.54%**

V. Discussion

To the best of our knowledge, this study is the first to address the segmentation of calf muscles using deep learning. The absence of prior work in this domain, coupled with our limited dataset size, presented a unique challenge. Therefore, our primary objective was to establish robust baselines and demonstrate a clear path toward improved performance under these constraints. Our systematic evaluation of seven models documented a progressive improvement in segmentation accuracy, providing a valuable methodological benchmark for this novel task.

The results culminate in SegModel_07 (our enhanced U-Net with data augmentation and fine-tuning), which significantly outperforms the initial baselines by achieving Dice scores of up to 79.85%. Data augmentation proved critical in preventing overfitting, as evidenced by the lower scores of non-augmented models (01-04). Furthermore, fine-tuning on all patient data improved performance, particularly for the PAD cohort, suggesting the model's robustness to pathological variations. The success of the enhanced U-Net architecture can be attributed to its design, which avoids excessive downsampling to preserve essential spatial details.

Despite these promising results, we acknowledge that the small dataset size (24 healthy subjects, 9 PAD patients) may limit generalizability. Future work will focus on validating our approach on larger, more diverse datasets, extending the methodology to other muscle groups, and exploring integration with clinical workflows for PAD diagnosis.

VI. Conclusions

We present a conditional GAN model with an enhanced U-Net for automated segmentation of lower limb calf muscles from DW-MRI, achieving high Dice score (54.86%-79.85%) from limited data through data augmentation and finetuning. This method offers a scalable solution for quantifying perfusion and diffusion particularly useful for post-treatment monitoring in PAD patients, reducing manual delineation effort and potentially enhancing the workflow for the automatic extraction of quantitative parameters characterizing micro- and macro-vascular flow in the lower limbs.

VII. Acknowledgment

We thank the radiographers at the Clinical Imaging Research Centre (CIRC), NUS, for the image acquisitions.

References

[1] M. R. Zemaitis, J. M. Boll, and M. A. Dreyer, "Peripheral Arterial Disease," in StatPearls [Internet]. Treasure Island (FL): StatPearls Publishing, Jan. 2025 [Updated May 23, 2023]. Available: https://www.ncbi.nlm.nih.gov/books/NBK430745/

[2] E. Pandey, "PMRI," GitHub repository, https://github.com/BesanHalwa/PMRI, accessed Aug. 30, 2025.

[3] K. A. Weber, R. Abbott, and V. Bojilov, "Multi-muscle segmentation of the cervical spine using convolutional neural networks on Dixon MRI," *NeuroImage: Clinical*, vol. 32, pp. 102–118, 2021.

[4] J. Ding, P. Cao, H. C. Chang *et al.*, "Deep learning-based thigh muscle segmentation for reproducible fat fraction quantification using fat–water decomposition MRI," *Insights into Imaging*, vol. 11, no. 1, p. 128, 2020, doi: 10.1186/s13244-020-00946-8.

[5] J. Zhu, B. Bolsterlee, B. V. Y. Chow, C. Cai, R. D. Herbert, Y. Song, and E. Meijering, "Deep learning methods for automatic segmentation of lower leg muscles and bones from MRI scans of children with and without cerebral palsy," *NMR in Biomedicine*, vol. 34, no. 12, p. e4609, Dec. 2021, doi: 10.1002/nbm.4609.

[6] A. Qureshi, S. Lim, S. Y. Suh, B. Mutawak, P. V. Chitnis, J. L. Demer, and Q. Wei, "Deep-learning-based segmentation of extraocular muscles from magnetic resonance images," Bioengineering (Basel), vol. 10, no. 6, p. 699, Jun. 2023, doi: 10.3390/bioengineering10060699.

[7] O. Ronneberger, P. Fischer, and T. Brox, "U-Net: Convolutional networks for biomedical image segmentation," in *Proc. Med. Image Comput. Comput.-Assist. Intervent. (MICCAI)*, 2015, vol. 9351, pp. 234–241, Springer, Cham, doi: 10.1007/978-3-319-24574-4_28.

[8] Yagis E, Aslani S, Jain Y, Zhou Y, Rahmani S, Brunet J, Bellier A, Werlein C, Ackermann M, Jonigk D, Tafforeau P, Lee PD, Walsh C. Deep Learning for 3D Vascular Segmentation in Phase Contrast Tomography. Res Sq [Preprint]. 2024 Jul 16:rs.3.rs-4613439. doi: 10.21203/rs.3.rs-4613439/v1. Update in: Sci Rep. 2024 Nov 8;14(1):27258. doi: 10.1038/s41598-024-77582-5. PMID: 39070623; PMCID: PMC11276017.

[9] Kim S, Hahn JO, Youn BD. Detection and Severity Assessment of Peripheral Occlusive Artery Disease via Deep Learning Analysis of Arterial Pulse Waveforms: Proof-of-Concept and Potential Challenges. Front Bioeng Biotechnol. 2020 Jun 30;8:720. doi: 10.3389/fbioe.2020.00720. PMID: 32714911; PMCID: PMC7340176.

[10] B. Khagi, T. Belousova, C. M. Short, A. A. Taylor, J. Bismuth, D. J. Shah, and G. Brunner, "Contrast-enhanced magnetic resonance imaging based calf muscle perfusion and machine learning in peripheral artery disease," *Scientific Reports*, vol. 15, no. 1, p. 4996, Feb. 2025, doi: 10.1038/s41598-025-87747-5.

[11] P. Isola, J. Y. Zhu, T. Zhou, and A. A. Efros, "Image-to-image translation with conditional adversarial networks," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, 2017, pp. 1125–1134.