International Research Journal on Advanced Engineering and Management
https://goldncloudpublications.com
https://doi.org/10.47392/IRJAEM.2026.0408
e ISSN: 2584-2854
Volume: 04 Issue: 10
October 2026
Page No: 2898 - 2902

EAT Deep Segmentation: Deep Learning-Based Automated Segmentation and Quantification of Epicardial Adipose Tissue from Cardiac CT Images

Dr. G. Saravanan[1], Ms. M. Amshavalli[2], R. S. Dhivya[3], M. Janani[4], G. Mohanraj[5]
[1,2,3,4,5]Department of Computer Science and Engineering, Erode Sengunthar Engineering College, Thudupathi, Erode – 638057, Tamil Nadu, India
Emails: *gsaravanan.pacet@gmail.com[1], amshavalli1494@gmail.com[2], rsdhivya25@gmail.com[3], jananimurugan513@gmail.com[4], mohanrpm2006rpm@gmail.com[5]*

Abstract
Epicardial adipose tissue (EAT) is the fat depot between the myocardium and the visceral pericardium, and its volume is an established marker of cardiometabolic risk. Measuring it on cardiac computed tomography (CT) requires manual slice-by-slice outlining, which is slow and varies between observers. This paper presents EAT Deep Segmentation, a deep learning framework for automated segmentation and quantification of EAT from cardiac CT images. The framework converts CT volumes and expert labels into paired two-dimensional slices, trains a U-Net encoder-decoder with skip connections to predict a pixel-wise EAT probability map, converts the map into a binary mask, and computes the EAT area from the mask and the pixel spacing. The model was implemented in TensorFlow and evaluated on a public EAT CT dataset, using 96 slices from 12 patients in a first experiment. Performance was measured with the Dice coefficient, intersection over union (IoU), precision and recall. The first training run gave a Dice coefficient of 0.277, an IoU of 0.172, a precision of 0.178 and a recall of 0.779. The high recall and low precision show that the network finds most EAT pixels but over-segments, which is expected with a small cohort and strong class imbalance. The work provides a reproducible baseline pipeline, and the analysis identifies larger patient cohorts, imbalance-aware loss functions and restriction to the pericardial region as the next steps toward clinically useful accuracy.
Keywords: *Cardiac CT; Deep learning; Epicardial adipose tissue; Medical image segmentation; U-Net*

1. Introduction
Epicardial adipose tissue (EAT) is the visceral fat depot located between the myocardium and the visceral layer of the pericardium. It shares its microcirculation with the myocardium and behaves as an active endocrine and paracrine organ, so changes in its volume are linked to coronary artery disease and metabolic risk (Iacobellis et al., 2005; Iacobellis, 2015). Population studies associate greater pericardial and epicardial fat with cardiovascular disease burden and incident myocardial infarction (Mahabadi et al., 2009, 2013). Cardiac computed tomography (CT) displays EAT as fat-density voxels, typically within a window of about −190 to −30 Hounsfield units (HU), inside the pericardial sac. Quantifying EAT therefore requires the fat inside the pericardium to be outlined slice by slice. Manual outlining is time consuming and depends on the observer, which limits its use in routine practice and in large studies. Early automated methods reduced this effort (Dey et al., 2008), and deep networks have since been used to quantify epicardial and thoracic fat from non-contrast CT (Commandeur et al., 2018). Segmenting EAT is difficult for three reasons. EAT forms a thin and irregular layer; its intensity is almost identical to that of adjacent paracardial fat; and it covers only a small fraction of each slice, so the foreground and background classes are strongly imbalanced. This paper presents EAT Deep Segmentation, an end-to-end pipeline that segments and quantifies EAT from cardiac CT images using a U-Net (Ronneberger et al., 2015).

1.1.Related work
Clinical reviews describe the local and systemic effects of EAT (Iacobellis, 2015) and its assessment by echocardiography (Iacobellis & Willens, 2009). On the algorithmic side, fully convolutional networks introduced end-to-end pixel-wise prediction (Shelhamer et al., 2017), and

International Research Journal on Advanced Engineering and Management
https://goldncloudpublications.com
https://doi.org/10.47392/IRJAEM.2026.0408
e ISSN: 2584-2854
Volume: 04 Issue: 10
October 2026
Page No: 2898 - 2902

U-Net added skip connections that work well with limited annotated data (Ronneberger et al., 2015). Extensions include 3D U-Net for volumetric data (Çiçek et al., 2016), V-Net with a Dice-based objective (Milletari et al., 2016) and the self-configuring nnU-Net (Isensee et al., 2021). A survey of these methods in medical image analysis is given by Litjens et al. (2017). Despite this progress, simple and reproducible EAT pipelines that can be trained on small public datasets are less commonly reported, which motivates the present work.

1.2.Contributions

The main contributions are: (i) a complete U-Net based pipeline from NRRD volumes and expert labels to a binary EAT mask and an EAT area measurement; (ii) an evaluation protocol based on the Dice coefficient, IoU, precision and recall; and (iii) a first-run baseline on a public dataset, together with an analysis of its main failure mode (over-segmentation) and a concrete improvement plan. Section 2 describes the method, Section 3 presents the results and discussion, and the final section concludes the paper.

2. Method

The proposed EAT Deep Segmentation framework has five stages, shown in Figure 1: data preparation, pre-processing, U-Net inference, thresholding, and quantification with evaluation. As shown in The experimental setup is summarised in Table 1. Figure 1 Proposed EAT Deep Segmentation pipeline.

Table 1 Experimental setup

Setting	Value
Dataset	Public figshare EAT cardiac CT dataset (NRRD)
Data used	96 axial slices from 12 patients
Task	Binary segmentation (EAT vs background)
Network	U-Net with skip connections
Framework	Python, TensorFlow

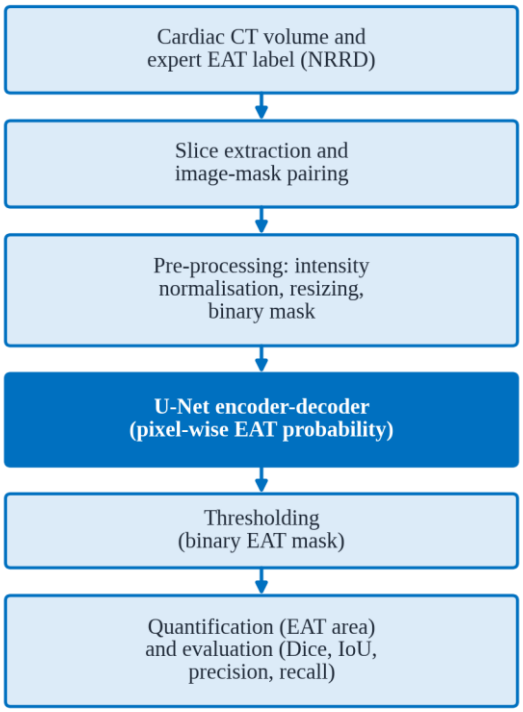


Figure 1 Proposed EAT Deep Segmentation pipeline

2.1.Dataset

The experiments use the publicly available epicardial adipose tissue cardiac CT dataset hosted on figshare (Siriapisith et al., 2021), in which the CT volumes and the expert EAT segmentations are stored in NRRD format. In the first experiment, 96 axial slices from 12 patients (the train_NC subset) were used.

2.2.Pre-processing

Each NRRD volume and its label volume were read and split into two-dimensional axial slices. Every slice was paired with the corresponding label slice, the CT intensities were normalised, the images were resized to a fixed input size, and the labels were binarised so that EAT pixels take the value 1 and all other pixels the value 0.

2.3.U-Net architecture

The segmentation network follows the U-Net design (Ronneberger et al., 2015), shown in Figure 2. The encoder repeats 3×3 convolutions followed by ReLU activations and 2×2 max-pooling, so the spatial size halves while the number of feature

International Research Journal on Advanced Engineering and Management
https://goldncloudpublications.com
https://doi.org/10.47392/IRJAEM.2026.0408
e ISSN: 2584-2854
Volume: 04 Issue: 10
October 2026
Page No: 2898 - 2902

channels increases. The bottleneck holds the most abstract context of the slice. The decoder restores the resolution with up-convolutions and concatenates the matching encoder feature maps through skip connections, which recovers fine boundaries. A final 1×1 convolution with a sigmoid activation gives a per-pixel EAT probability p(x, y). The binary mask is obtained by thresholding:

$$M(x, y) = 1 \text{ if } p(x, y) \geq \tau, \text{ otherwise } 0 \quad (1)$$

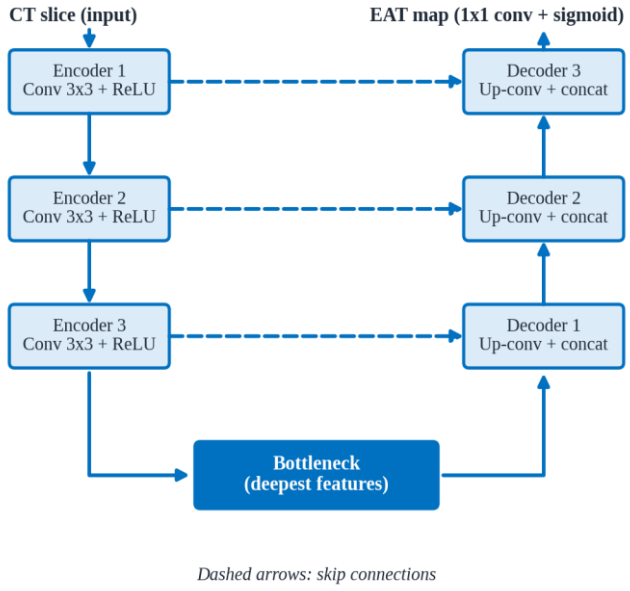


Figure 2 U-Net architecture used for EAT segmentation (dashed arrows: skip connections)

2.4.Training and implementation

The network was implemented in Python using TensorFlow (Abadi et al., 2016) and trained on a Windows laptop. The dataset preparation script converts the volumes into slice and mask arrays, and the training script fits the network and evaluates it on the test slices. The experimental setup is summarised in Table 1 [1 – 10].

2.5.Quantification

The EAT area of a slice is computed from the predicted mask and the in-plane pixel spacing stored in the NRRD header:

$$A = N \times \Delta x \times \Delta y \quad (2)$$

where N is the number of EAT pixels in the mask and Δx, Δy are the pixel spacings in mm, so A is given in mm². For a set of consecutive slices with thickness t, the EAT volume is

$$V = t \times \Sigma A_k \quad (3)$$

where Ak is the area of slice k.

2.6.Evaluation metrics

Segmentation quality was measured at pixel level by comparing the predicted mask with the expert mask (Taha & Hanbury, 2015). TP, FP and FN denote the numbers of correctly marked EAT pixels, pixels marked as EAT that are not EAT, and expert EAT pixels that were missed. The four metrics are defined As shown in Table 2 Evaluation metrics.

Table 2 Evaluation metrics

Metric	Definition
Dice	2TP / (2TP + FP + FN)
IoU	TP / (TP + FP + FN)
Precision	TP / (TP + FP)
Recall	TP / (TP + FN)

3. Results And Discussion

3.1.Results

On the test set, the first training run gave a Dice coefficient of 0.277, an IoU of 0.172, a precision of 0.178 and a recall of 0.779 (Table 3, Figure 3). The recall is about 4.4 times the precision: roughly 78% of the expert-labelled EAT pixels were recovered, but only about 18% of the predicted EAT pixels agree with the expert label. As shown in Table 3 Test results of the first training run.

Table 3 Test results of the first training run

Metric	Test score
Dice	0.277
IoU	0.172
Precision	0.178

International Research Journal on Advanced Engineering and Management
https://goldncloudpublications.com
https://doi.org/10.47392/IRJAEM.2026.0408
e ISSN: 2584-2854
Volume: 04 Issue: 10
October 2026
Page No: 2898 - 2902

Recall	0.779

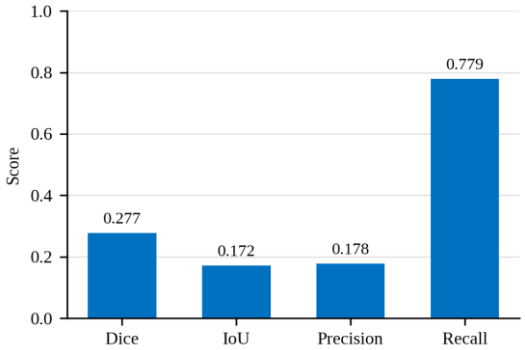


Figure 3 Test metrics of the first training run

3.2.Discussion

High recall with low precision means that false positives clearly outnumber false negatives: the network rarely misses EAT, but it marks a much larger region than the expert label. Dice and IoU penalise both error types, so they remain low because of the false positives even though the sensitivity is good. Three factors probably explain this behaviour. First, only 12 patients were used, so the network sees limited anatomical variation. Second, EAT is a small foreground class, which makes the loss function and the decision threshold critical (Lin et al., 2017; Sudre et al., 2017). Third, EAT has the same HU range as paracardial and mediastinal fat outside the pericardium. Without an explicit constraint to the pericardial region, a slice-wise network can label such fat as EAT, which would produce false positives of the kind observed here. The study has clear limitations: a small cohort, two-dimensional slice-wise segmentation, a single training run without cross-validation or confidence intervals, and no pericardial post-processing. The reported figures are therefore a baseline and not clinical accuracy, and prior work on larger cohorts reports considerably higher agreement with experts [11 – 17] (Commandeur et al., 2018). The planned improvements are a larger patient cohort, data augmentation (Shorten & Khoshgoftaar, 2019), imbalance-aware losses such as Dice and focal loss, restriction of predictions to the pericardial region with morphological post-processing, and volumetric models (Çiçek et al., 2016; Isensee et al., 2021).

Conclusion

This paper presented EAT Deep Segmentation, an end-to-end U-Net pipeline that converts cardiac CT volumes and expert labels into a binary EAT mask and an EAT area measurement. On a first experiment with 96 slices from 12 patients, the model reached a Dice coefficient of 0.277, an IoU of 0.172, a precision of 0.178 and a recall of 0.779. The network detects most EAT pixels but over-segments, so the result is a reproducible baseline rather than a final accuracy claim. Future work will use more patients, augmentation, imbalance-aware losses, pericardial region restriction and three-dimensional models to reach accuracy suitable for clinical use.

Acknowledgements

The authors thank the Department of Computer Science and Engineering, Erode Sengunthar Engineering College, for its support, and the providers of the public EAT cardiac CT dataset.

References

[1]. Iacobellis, G., Corradi, D., & Sharma, A. M. (2005). Epicardial adipose tissue: Anatomic, biomolecular and clinical relationships with the heart. Nature Clinical Practice Cardiovascular Medicine, 2(10), 536–543.

[2]. Iacobellis, G. (2015). Local and systemic effects of the multifaceted epicardial adipose tissue depot. Nature Reviews Endocrinology, 11(6), 363–371.

[3]. Mahabadi, A. A., et al. (2009). Association of pericardial fat, intrathoracic fat, and visceral abdominal fat with cardiovascular disease burden: The Framingham Heart Study. European Heart Journal, 30(7), 850–856.

[4]. Mahabadi, A. A., et al. (2013). Association of epicardial fat with cardiovascular risk factors and incident myocardial infarction in the general population: The Heinz Nixdorf Recall Study. Journal of the American College of Cardiology, 61(13), 1388–1395.

[5]. Dey, D., et al. (2008). Automated

International Research Journal on Advanced Engineering and Management
https://goldncloudpublications.com
https://doi.org/10.47392/IRJAEM.2026.0408
e ISSN: 2584-2854
Volume: 04 Issue: 10
October 2026
Page No: 2898 - 2902

quantitation of pericardiac fat from noncontrast CT. Investigative Radiology, 43(2), 145–153.

[6]. Commandeur, F., et al. (2018). Deep learning for quantification of epicardial and thoracic adipose tissue from non-contrast CT. IEEE Transactions on Medical Imaging, 37(8), 1835–1846.

[7]. Ronneberger, O., Fischer, P., & Brox, T. (2015). U-Net: Convolutional networks for biomedical image segmentation. In Medical Image Computing and Computer-Assisted Intervention – MICCAI 2015 (LNCS 9351, pp. 234–241). Springer.

[8]. Iacobellis, G., & Willens, H. J. (2009). Echocardiographic epicardial fat: A review of research and clinical applications. Journal of the American Society of Echocardiography, 22(12), 1311–1319.

[9]. Shelhamer, E., Long, J., & Darrell, T. (2017). Fully convolutional networks for semantic segmentation. IEEE Transactions on Pattern Analysis and Machine Intelligence, 39(4), 640–651.

[10]. Çiçek, Ö., Abdulkadir, A., Lienkamp, S. S., Brox, T., & Ronneberger, O. (2016). 3D U-Net: Learning dense volumetric segmentation from sparse annotation. In MICCAI 2016 (LNCS 9901, pp. 424–432). Springer.

[11]. Milletari, F., Navab, N., & Ahmadi, S.-A. (2016). V-Net: Fully convolutional neural networks for volumetric medical image segmentation. In Proceedings of the Fourth International Conference on 3D Vision (3DV) (pp. 565–571). IEEE.

[12]. Isensee, F., Jaeger, P. F., Kohl, S. A. A., Petersen, J., & Maier-Hein, K. H. (2021). nnU-Net: A self-configuring method for deep learning-based biomedical image segmentation. Nature Methods, 18(2), 203–211.

[13]. Litjens, G., Kooi, T., Bejnordi, B. E., Setio, A. A. A., Ciompi, F., Ghafoorian, M., van der Laak, J. A. W. M., van Ginneken, B., & Sánchez, C. I. (2017). A survey on deep learning in medical image analysis. Medical Image Analysis, 42, 60–88.

[14]. Siriapisith, T., et al. (2021). Epicardial adipose tissue cardiac CT dataset [Data set]. figshare.

[15]. Abadi, M., Barham, P., Chen, J., Chen, Z., Davis, A., Dean, J., … Zheng, X. (2016). TensorFlow: A system for large-scale machine learning. In Proceedings of the 12th USENIX Symposium on Operating Systems Design and Implementation (OSDI) (pp. 265–283). USENIX Association.

[16]. Taha, A. A., & Hanbury, A. (2015). Metrics for evaluating 3D medical image segmentation: Analysis, selection, and tool. BMC Medical Imaging, 15, Article 29.

[17]. Lin, T.-Y., Goyal, P., Girshick, R., He, K., & Dollár, P. (2017). Focal loss for dense object detection. In Proceedings of the IEEE International Conference on Computer Vision (ICCV) (pp. 2980–2988). IEEE.

[18]. Sudre, C. H., Li, W., Vercauteren, T., Ourselin, S., & Cardoso, M. J. (2017). Generalised Dice overlap as a deep learning loss function for highly unbalanced segmentations. In Deep Learning in Medical Image Analysis and Multimodal Learning for Clinical Decision Support (LNCS 10553, pp. 240–248). Springer.

[19]. Shorten, C., & Khoshgoftaar, T. M. (2019). A survey on image data augmentation for deep learning. Journal of Big Data, 6, Article 60.