Decoding Stroke Patterns: A Novel Deep Learning Approach to Atrial Fibrillation Risk Stratification

Mohammad Javad Shokri*, Nandakishor Desai*, Aravinda S. Rao*, Angelos Sharobeam†, Bernard Yan†, and Marimuthu Palaniswami*
* Electrical and Electronic Engineering, The University of Melbourne, Parkville, Victoria - 3010, Australia
Email: mshokri@student.unimelb.edu.au; {nnandakisho, aravinda.rao, palani}@unimelb.edu.au
† Melbourne Brain Center, Royal Melbourne Hospital, Parkville, Victoria - 3052, Australia
Email: angelos.sharobeam@mh.org.au, bernard.yan@mh.org.au

***Abstract*—Atrial fibrillation (AF), being the most prevalent arrhythmia around the world, is a significant health concern considering an aging population and increasing prevalence of its risk factors such as hypertension and obesity. It is estimated that AF increases the risk of stroke by about five times and the risk of its recurrence by two-fold. AF remains undetected in up to 30% of cases due to its asymptomatic and paroxysmal nature, and lack of routine screening. We present a novel AF risk stratification framework using brain magnetic resonance imaging (MRI) to identify the underlying AF in post-stroke patients and assist in preventing secondary ones. By analyzing the infarct patterns of these patients using a multi-task learning framework (adopting segmentation and classification losses simultaneously), our proposed model achieves an area under the receiver operating characteristic (AUROC) of 87.48 $\pm$ 4.88, demonstrating its capability in discriminating AF patients from others. Since MRI is already an established modality in the stroke treatment and diagnosis framework, this innovative solution incurs no additional costs or tests for the patient. It can effectively select patients at elevated risk for extensive cardiac investigation and definite diagnosis of AF.**

I. Introduction

Atrial fibrillation (AF) is an irregular and often rapid heart rhythm that can lead to blood clots in the heart. A blood clot detached from the heart can travel to the brain and result in an ischemic stroke. AF is a significant risk factor for stroke, responsible for up to one-third of all cases [14]. AF, characterized by a chaotic and irregular heart rhythm, increases stroke risk up to five times [22]. If left untreated, AF can raise the risk of recurrent strokes by two times [21], making early diagnosis and management essential. Detecting AF soon after a stroke is vital, as timely administration of anticoagulation therapy can significantly reduce the risk of subsequent strokes by helping prevent the formation of blood clots. However, up to 30% of AF cases remain undetected, primarily due to its intermittent and often asymptomatic nature. Many individuals with AF experience sporadic episodes that may go unnoticed during routine or even prolonged monitoring periods. This intermittent nature complicates detection, as traditional diagnostic methods like electrocardiograms (ECGs) or Holter monitors may fail to capture these transient episodes. Additionally, AF can appear with subtle or no symptoms, further delaying diagnosis. The lack of observable symptoms also makes it challenging for patients to recognize the need for medical evaluation and for healthcare providers to initiate timely diagnostic procedures. Furthermore, the diagnostic tools available for AF detection, such as prolonged ECG monitoring devices and implantable cardiac monitors, can be costly, invasive, and require significant patient compliance. These factors contribute to the challenges in achieving widespread and efficient AF screening, especially in populations with limited access to healthcare resources. These challenges underscore the need for an innovative, cost-effective approach to improve AF detection and management. Due to the impracticality of extensive cardiac monitoring for all post-stroke patients, comprehensive risk assessment is critical to select the patients at high risk for further investigation.

Recently, smartphones and wearable devices with heart rate sensors have enabled the monitoring of heart rhythms outside traditional clinical settings [31]. These devices streamline AF detection through continuous and convenient monitoring but face challenges, including usability barriers for older adults, high costs, and integration issues with clinical workflows. On the other hand, brain imaging is crucial for stroke diagnosis and management, with various modalities available in stroke centers. Computed tomography (CT) scans are typically the first choice in acute stroke settings due to their rapid availability and ability to quickly differentiate between ischemic and hemorrhagic strokes. However, magnetic resonance imaging (MRI)—particularly diffusion-weighted imaging (DWI)—has also become routine soon after a stroke. Despite its longer scanning duration, it is essential in the acute phase of stroke as it offers highly detailed images of brain tissue with superior sensitivity and specificity. MRI can detect small or unusual lesions and differentiate between ischemic and hemorrhagic strokes, which is vital for treatment planning. While wearable technology has its benefits, routinely performed MRIs can also be leveraged to diagnose AF soon after a stroke. Our proposed deep learning solution utilizes existing MRIs to detect underlying AF, offering an effective and cost-efficient method that integrates seamlessly into current stroke diagnosis procedures while requiring no additional expenses or examinations for the patient.

Research into AF patterns in brain MRI using machine learning is limited. Previous works have analyzed stroke etiology, infarct topography, and volume in cardioembolic strokes, focusing on AF strokes [24]. Studies such as [19] have examined the association of AF with specific white matter hyperintensity patterns in embolic stroke patients. In contrast, [9] highlighted the burden of silent brain lesions in AF patients, suggesting the utility of MRI screening. The work in [15] explored brain MRI for personalized AF therapy, and [11] integrated MRI and blood markers for stroke risk insights. Further, [6] assessed the impact of including silent brain infarcts from MRI in the CHA2DS2-VASc score for stroke risk management in AF patients.

Despite advancements, distinguishing AF-related strokes solely from brain images remains a clinical challenge due to limited knowledge of AF-associated infarct patterns. Furthermore, the limited works focus primarily on explorative analysis without exploiting the benefits of deep learning paradigms to develop high-performing automated workflows to detect the presence of AF. We suggest that deep learning (DL) models, which excel in automatic feature extraction, can enable the identification of AF through stroke imaging—a task previously not possible. DL models are particularly suitable for medical imaging due to established digital workflows and standards in radiology image storage. To achieve AF identification, we propose a multi-task learning approach by integrating segmentation and classification of brain MRI images. By simultaneously learning to delineate regions of interest (infarcts) and classify these images, our method aims to enhance both accuracy and interpretability in AF risk assessment. To the best of our knowledge, our proposed approach is the first such technique designed for early detection of AF from brain MRI scans. The key contributions of this work include:

- Introducing an innovative approach to integrate AF diagnosis into routine stroke imaging protocols, which can aid in more accurate treatment and management strategies.
- Developing a novel deep learning framework to differentiate AF from other stroke etiologies in post-stroke brain MRI.
- Employing a multi-task learning approach that combines segmentation and classification tasks, improving both model performance and interpretability.

II. Related Works

DL algorithms have been widely adopted across various applications in stroke imaging, including diagnosis, detection, risk prediction, prognosis, treatment recommendation, medication management, and patient monitoring [13], [7], [23], [25]. Some studies specifically focus on stroke MRI imaging, addressing tasks such as time since stroke onset (TSS) classification [16], prediction of final infarct region and tissue outcomes [27], clinical outcomes prediction [12], [10], and segmentation and detection of abnormalities like infarcts [29], [18].

Despite the rapid growth of this field, to the best of our knowledge, only one concurrent study in the literature has explored a similar approach to ours. It is important to note that our research was conducted independently and was not influenced by this work, as it was not published during the initial stages of our investigation. Zhang *et al*. employed a convolutional neural network (CNN) to extract features from brain MRI, which were subsequently combined with radiomic features from segmented images to develop and validate a classification algorithm [30]. Their model achieved an area under the receiver operating characteristic curve (AUC-ROC) of 79.9%, showcasing the potential of utilizing brain MRI for detecting underlying AF. However, details regarding the control group of the dataset, network architecture, and specifics of the radiomic features extracted were not provided in their study. Moreover, the absence of comprehensive experiments and discussion limits the assessment of its distinct contributions.

III. Methodology

A. Neuroimaging for AF Screening

DWI-MRI is routinely performed on many stroke patients in stroke centers. As a result, an automated MRI-based risk stratification framework could be pivotal for the early identification of AF patients who would benefit most from prolonged cardiac monitoring. By utilizing a deep learning model to analyze brain MRI images, the likelihood of AF can be estimated based on infarct patterns. This estimation then serves as a risk stratification tool, identifying high-risk patients who should undergo extensive cardiac monitoring to definitively diagnose AF. To develop this framework, we designed a model that leverages a multi-task approach, combining classification and segmentation tasks. This method ensures that the classification network specifically targets and analyzes the features of the infarct regions.

B. Dataset

To evaluate the proposed approach for AF identification, a dataset consisting of 235 acute ischemic stroke (AIS) patients was retrospectively acquired. Ethics approval was obtained before the commencement of the project from the Royal Melbourne Hospital Ethics Committee (QA 2013.072). For each patient, the data includes a 3D brain DWI-MRI image, a corresponding class label for the AF or control group, and a segmentation mask highlighting the infarct regions within the brain. The MRI images were acquired at patient admission during the acute phase of the stroke. The imaging was performed using various scanners from Siemens Aera, Siemens Prisma Fit, Siemens Skyra, Siemens Magnetom Essenza, and Philips Ingenia. The dataset demographics and imaging parameters are illustrated in Table I. The annotation process for this dataset involves two parts of class labeling and segmentation masks for the infarct regions, described in the following:

To identify the underlying cause of stroke based on its etiology, different systems have been developed including the Causative Classification System (CCS) [4], Trial of Org10172 in Acute Stroke Treatment (TOAST) [1], and ASCO (atherosclerosis, small-vessel disease, cardiac source, and other cause) [2]. These systems provide frameworks for

TABLE I
OVERVIEW OF DATASET DEMOGRAPHICS AND IMAGING PARAMETERS, INCLUDING PATIENT CHARACTERISTICS, IMAGING PROTOCOLS, AND PARAMETERS

Number of patients	235
AF-related strokes	138 (58.7%)
Female	83 (35.3%)
Age (mean±std)	71.1 ± 14.2
Magnetic field strength	1.5 or 3 Tesla
Repetition time	4100-7920 ms
Echo time	55-104 ms
Flip angle	0-180 degrees
b-values	0 or 1,000 sec/mm^2
Slice thickness	0.256-7.474 mm
Slice spacing	2-7.5 mm
Pixel spacing	0.548-2 mm

categorizing strokes based on underlying causes, aiding in diagnosis and treatment planning. Both TOAST and CCS assess clinical features, imaging findings, and other diagnostic criteria to categorize strokes into specific etiological groups. These categories include large artery atherosclerosis (LAA), cardioembolism (CE), small artery occlusion (SAO), stroke of other determined cause (OC), and stroke of undetermined cause (UND). Among these systems, CCS offers a more precise categorization of stroke causes, with more significant inter-category variability than intra-category variability. CCS reassigns 20-40% of cases from the undetermined category in other systems to specific subtypes, providing enhanced discrimination for clinical, imaging, and prognostic characteristics. The unknown category is markedly smaller in CCS (33%) compared to TOAST (53%) and ASCO (42%), highlighting its superior accuracy in categorizing stroke etiologies [3].

In this study, two expert neurologists utilized (CCS) to determine whether the underlying stroke etiology is AF, a significant subgroup of cardioembolic (CE) strokes. Patients are categorized into two groups: (1) strokes associated with AF and (2) strokes caused by LAA. AF-related strokes are identified through clinical reports or visual analysis of a 12-lead ECG, excluding newly diagnosed AF within 30 days post-cardiac surgery and strokes resulting from other cardioembolic sources. In contrast, LAA-related strokes involve infarctions associated with significant stenosis of the parent artery, with other potential cardioembolic sources reasonably excluded based on CCS criteria. Furthermore, the neurologists utilized the ITK-SNAP software [28] to annotate the DWI-MRI images, generating segmentation ground truth masks for this study.

C. Network Architecture

The proposed model, illustrated in Figure 1, integrates a pre-trained ResNet-18 with a custom decoder and a classification head, tailored for automated stroke etiology classification in brain MRI scans. The details of these components are as follows:

D. Encoder (Backbone)

A 3D ResNet-18 is selected as the encoder due to its proven effectiveness in various medical image analysis tasks, particularly when dealing with small datasets. We use a pre-trained variant of this model from [8], originally developed as a part of a medical image segmentation network. Specifically adapted for 3D volumetric data, the encoder processes the single-channel brain MRI scans to capture abstract features through residual blocks.

This encoded output can be fed into a decoder for a segmentation task, which aims to delineate infarct areas in brain MRIs. Alternatively, the encoded features can be utilized for classification tasks, distinguishing between different types of strokes. In addition, the model architecture supports a multi-task learning approach, where both segmentation and classification are performed simultaneously.

E. Decoder

The decoder in our framework complements the encoder and classification head by reconstructing spatial representations crucial for segmenting infarct regions from brain MRI images. It utilizes transposed convolutional layers to upsample encoded feature maps, restoring them to their original dimensions. This process facilitates the localization of infarcted areas based on learned features, integrating information across scales captured during encoding. By working in tandem with the classification head (which identifies stroke class), the decoder enhances interpretability, providing insights into spatial distribution and extent of infarcts.

F. Classification Head

The classification head in our framework plays a crucial role in categorizing stroke subtypes, particularly focusing on AF, from brain MRI images. Building upon the encoded features from the encoder, the classification head utilizes a global average pooling layer (GAP) and a linear layer to perform classification. It integrates information across the encoded representation to predict the presence of AF-related patterns within the MRI scans.

By leveraging deep learning techniques, the classification head enhances the model's ability to discern subtle variations indicative of AF, contributing to more accurate diagnosis and risk assessment. This component not only identifies AF but also provides probabilistic scores, aiding in the clinical interpretation of findings and prioritizing patients according to their risk levels.

IV. EXPERIMENTS

To evaluate our proposed framework, we carry out comprehensive experiments. First, we compare our model with the existing work in the literature. Next, we conduct ablation studies to assess the impact of the multi-task learning approach. Further, we visually inspect the model's activation maps to investigate its decision-making process.

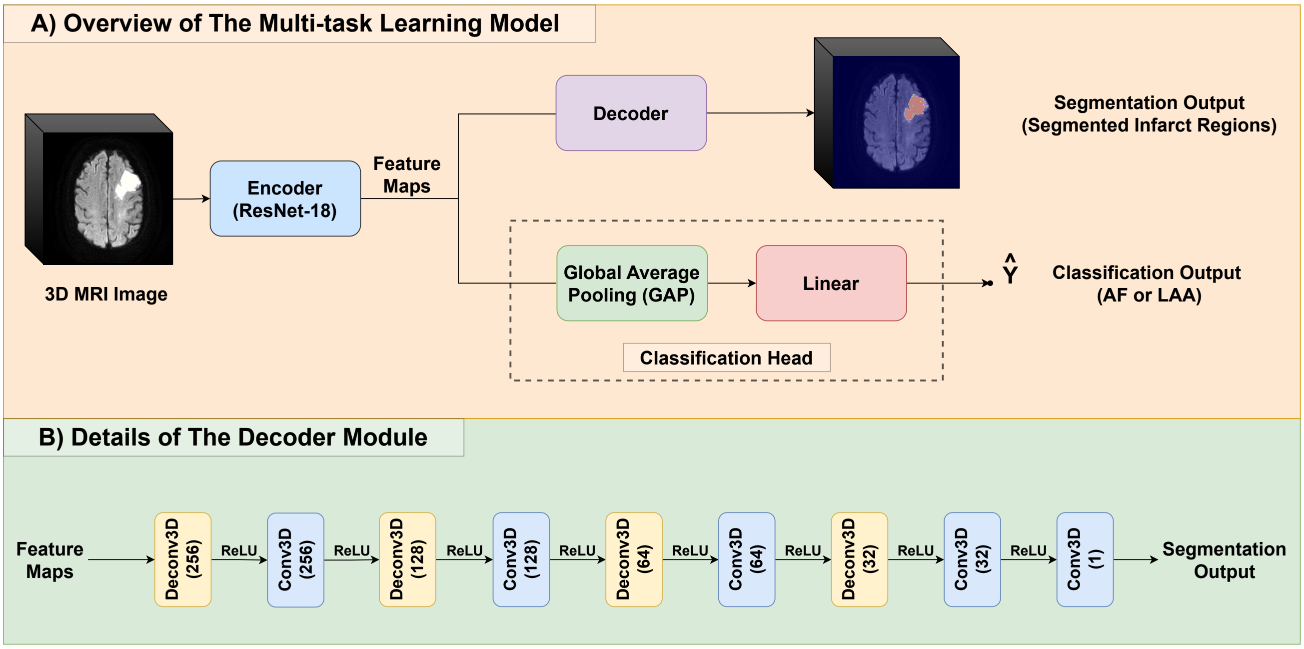

Fig. 1. Architecture of the proposed multi-task learning network for AF identification from brain MRI. (A) Overview of the model showing the input MRI image processed through the encoder (pre-trained 3D ResNet-18) to extract features, followed by a decoder and a classification head in parallel. The model outputs include segmentation masks and predicted class labels. (B) Detailed view of the decoder module highlighting 3D deconvolution layers that upsample feature maps to reconstruct segmentation details and 3D convolution layers to refine the upsampled features. The number of output channels in each layer is shown in parentheses.

A. Pre-processing and Data Augmentation

Pre-processing of medical images plays a critical role in their analysis. In this study, we utilized simple yet effective steps to ensure consistent spatial and physical resolutions throughout the dataset. Standardizing these resolutions enhances training efficacy, accelerates convergence, and improves overall performance. Specifically, all images and their corresponding segmentation masks underwent resampling using the nearest interpolation method to achieve a physical resolution of (8mm, 1mm, 1mm) in the z, x, and y directions, respectively. Furthermore, both the images and masks were resized to (32, 256, 256) dimensions. In addition, intensity values were clipped to fall within the range of the 0.05$^{\text{th}}$ and 99.5$^{\text{th}}$ percentiles, and normalization was applied using the min-max approach across all samples.

We also applied a series of data augmentations using the Volumentations 3D package [26] to enhance robustness and generalization in our training data. These included elastic transformations, independent flips along the three axes, random 90-degree rotation in the axial plane, Gaussian noise addition, random cropping, and scaling.

B. Loss Functions

To train our model, we employ the summation of two different loss functions tailored for our specific tasks: binary cross-entropy with logits loss (BCEWithLogitsLoss) for the AF identification task and a combination of Dice loss and cross-entropy loss (DiceCELoss) for the segmentation task. The binary cross-entropy with logits loss, denoted as $\mathcal{L}_{\text{BCE}}$, is used for the AF identification task. This loss function combines a sigmoid layer and the binary cross-entropy loss in one single class. It is defined as:

$$\mathcal{L}_{\text{BCE}} = -\frac{1}{N}\sum_{i=1}^{N}\left[y_i \log \hat{y}_i + (1-y_i)\log(1-\hat{y}_i)\right] \quad (1)$$

where N is the number of samples, y_i is the ground truth label (AF or LAA) for the i^{th} sample, $\hat{y}_i$ is the predicted probability for the i^{th} sample, obtained by applying the sigmoid function (σ) to the raw prediction (logit) x_i, *i.e.*,,

$$\hat{y}_i = \sigma(x_i) = \frac{1}{1+e^{-x_i}}. \quad (2)$$

For the segmentation task, we use a weighted combination of Dice loss and cross-entropy loss. The Dice coefficient is particularly effective for handling imbalanced data by focusing on the overlap between the predicted and ground truth masks. The cross-entropy loss focuses on the classification error for each voxel within the image. The combined Dice cross-entropy loss, denoted as $\mathcal{L}_{\text{DiceCE}}$, is defined as:

$$\mathcal{L}_{\text{DiceCE}} = \lambda_{\text{dice}}\mathcal{L}_{\text{Dice}} + \lambda_{\text{ce}}\mathcal{L}_{\text{CE}} \quad (3)$$

where λ_{dice} and λ_{ce} are the weights for the Dice loss and cross-entropy loss components, respectively. The Dice loss, $\mathcal{L}_{\text{Dice}}$, is given by:

$$\mathcal{L}_{\text{Dice}} = 1 - \frac{1}{N}\sum_{n=1}^{N}\frac{2\sum_{i=1}^{V} p_{ni}g_{ni}}{\sum_{i=1}^{V}\left[p_{ni}+g_{ni}\right]} \quad (4)$$

TABLE II

COMPARISON WITH THE OTHER METHOD DEMONSTRATES THE EFFICACY OF OUR METHODOLOGY, SIGNIFICANTLY SURPASSING PREVIOUS RESEARCH.

Method	Accuracy	Precision	Recall	F1 score	AUC-ROC
Zhang et al. [30]	70	63.8	92.5	75.5	79.9
Ours	**80.87 ± 5.57**	**89.72 ± 5.50**	78.10 ± 13.75	**82.48 ± 7.51**	**87.48 ± 4.88**

TABLE III

ABLATION STUDIES: WE PERFORMED AF IDENTIFICATION USING THE ENCODER AND THE CLASSIFICATION HEAD, UTILIZING THE CLASSIFICATION LOSS SOLELY AND OMITTING THE SEGMENTATION MODULE AND LOSS. THE RESULTS UNDERSCORE THE EFFICACY OF THE MULTI-TASK LEARNING APPROACH, SIGNIFICANTLY OUTPERFORMING THE CLASSIFICATION-ONLY METHOD.

Method	Accuracy	Precision	Recall	F1 score	AUC-ROC
ResNet-18 (Classification)	73.48 ± 4.64	86.84 ± 5.67	62.55 ± 9.48	72.41 ± 7.74	80.25 ± 3.20
ResNet-18 (Multi-Task)	**80.87 ± 5.57**	**89.72 ± 5.50**	**78.10 ± 13.75**	**82.48 ± 7.51**	**87.48 ± 4.88**

where N is the number of samples, V is the number of voxels in each sample, p_{ni} is the predicted probability for voxel i in sample n, and g_{ni} is the ground truth label for voxel i in sample n. And similar to (1), $\mathcal{L}_{\text{CE}}$ for the segmentation task for one image is given by:

$$\mathcal{L}_{\text{CE}} = -\frac{1}{V}\sum_{i=1}^{V}\left[g_i \log p_i + (1-g_i)\log(1-p_i)\right] \quad (5)$$

Where p_i represents the predicted probability for the i^{th} voxel, g_i is the ground truth label for the i^{th} voxel, and V denotes the total number of voxels. Subsequently, this loss is averaged across all samples.

We hypothesize that including the segmentation loss in the training process could potentially enhance the classification performance, particularly for stroke etiologies and AF identification, by forcing the model to focus on the infarct regions. By learning to segment these areas, the model gains a better understanding of the spatial and morphological features associated with each etiology, thereby improving its ability to classify AF from brain MRI scans accurately.

C. Implementation Details

The models were created using PyTorch [20] and trained on an NVIDIA A100 GPU. Stochastic gradient descent (SGD) served as the optimizer (learning rate for the encoder and decoder = 0.0001, learning rate for the classification head = 0.001, decay rate for step scheduler = 0.5, and scheduler step size = 30). Training was performed for 200 epochs, with a batch size set to 4. To mitigate overfitting, L_2 regularization with a weight decay of 0.001 was applied. In our implementation, we set $\lambda_{\text{dice}} = 0.3$ and $\lambda_{\text{ce}} = 0.7$, balancing the contributions of each loss component to the overall segmentation loss.

V. EVALUATION

The model underwent rigorous evaluation via five-fold cross-validation. This method involved partitioning the dataset into five subsets, each containing 20% of the total images, employing stratified sampling to maintain balanced class distributions across splits. Subsequently, the network employed five independent training and evaluation cycles. During each iteration, four subsets were utilized for training, while one subset was used for evaluating the model's performance. This approach ensured a robust assessment of the model's generalization and effectiveness across diverse subsets of the dataset. We evaluated the performance of our classification framework using standard metrics, including accuracy, precision, recall (sensitivity), F1 score, and area under the receiver operating characteristic curve (AUC-ROC).

A. AF Identification

We compare our AF identification model with the previous work in the literature. [30] develop a model using MRI to detect AF in stroke patients. They focused on combining radiomic and semantic features extracted from CNNs to achieve this. They used a dataset consisting of 489 patients to validate their model, achieving an AUC-ROC of 79.9%. Through this work, they highlighted a potential link between ischemic lesion patterns and AF etiology. However, the lack of comprehensive experiments and discussion makes it difficult to assess its contributions.

Instead of extracting features from segmented images similar to [30], we employed a multi-task learning approach to train a segmentation and classification network simultaneously to identify AF. This approach involved leveraging shared features between segmentation and classification tasks. Our approach yielded results that significantly outperformed theirs, as demonstrated in Table II. Our model attained an AUC-ROC of 87.48 ± 4.88, demonstrating that integrating segmentation with classification tasks not only improves AF identification accuracy but also enhances the interpretability of feature extraction related to infarct regions.

B. Ablation Study

To investigate the effectiveness of the multi-task learning approach and analyze its impact, we conducted ablation studies. In these experiments, we evaluated AF identification using the encoder and the classification head, utilizing solely

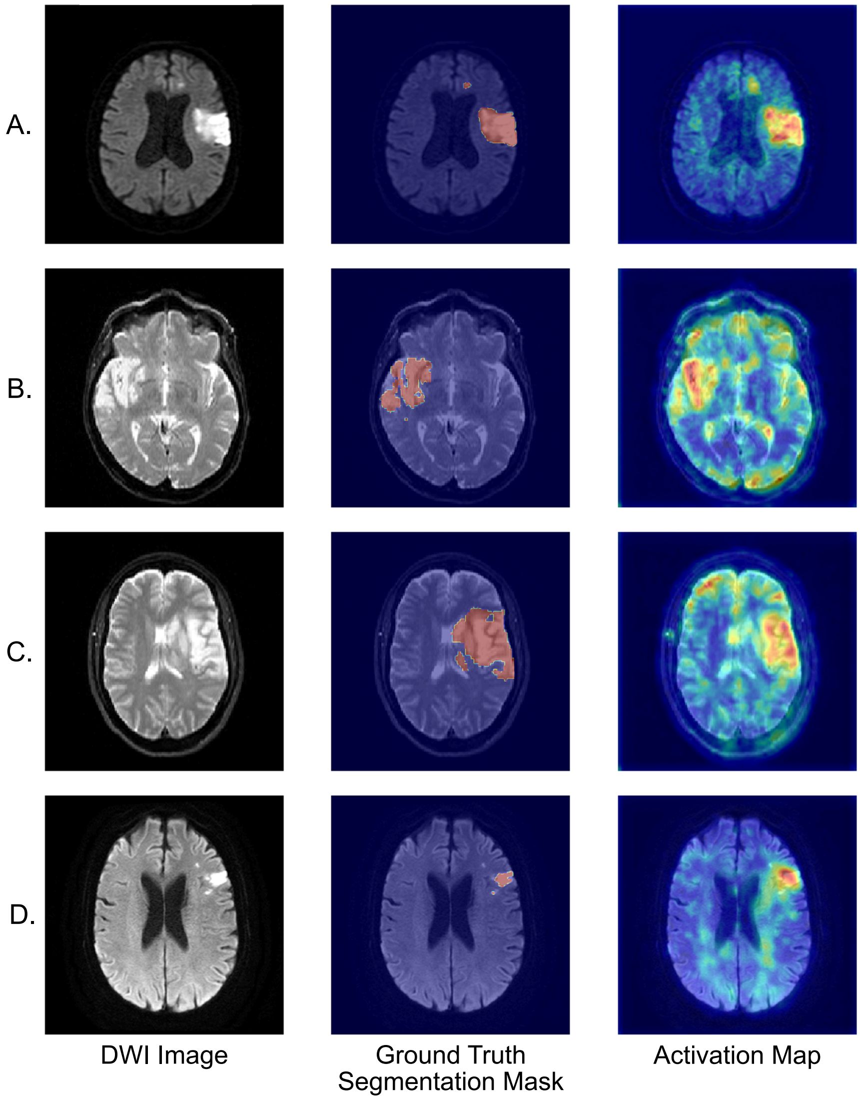

Fig. 2. Visual analysis of the proposed method's performance in distinguishing between AF (A, B) and LAA (C, D) groups: From left to right, the first image is an axial slice of the DWI images. The second image displays the ground truth segmentation masks, highlighting the infarct region within the image. The highlighted areas in red indicate regions identified by experts as infarcted tissue. The third image is the activation map of the encoder's last layer. The activation maps indicate that, in the adopted multi-task learning approach, the model focuses mainly on infarct areas for decision-making regarding stroke etiology identification.

the classification loss and omitting the segmentation module. The results, detailed in Table III, underscore the efficacy of our approach. The multi-task learning approach significantly outperforms the classification-only method, highlighting its effectiveness. This superiority can be attributed to the use of the segmentation masks and decoder, along with the segmentation loss, which compels the encoder to extract features mainly related to the infarct regions. These features, which are also shared with the classification head, enhance the performance of AF identification.

C. *Visual Interpretation*

To visually interpret the behavior of our model, we evaluated the activation maps generated from the last layer in the encoder for both classes of AF and LAA. Samples of these maps, along with an axial slice of the DWI images and ground truth segmentation masks highlighting the location of the infarct regions, are shown in Figure 2. The activation maps show areas of the brain that the model considers important for its decision-making process. In our multi-task learning approach (involving both classification and segmentation tasks), these activation maps reveal that the model focuses mainly on the infarct regions. This suggests that the model leverages these areas to make informed decisions about stroke etiology, distinguishing between AF and LAA patients.

D. Discussion

Our research aimed to identify AF from brain MRI scans in post-stroke patients to assist in preventing AF-related secondary strokes. To achieve this, we adopted a multi-task learning approach emphasizing the encoding of infarct regions by utilizing a segmentation loss. This strategy was designed to enhance the classification performance of our model, which we successfully demonstrated by achieving state-of-the-art results. Our approach significantly outperformed the previous work, showcasing its efficacy in AF identification.

It is important to note that our model did not achieve satisfactory performance in segmentation. This limitation can be attributed to several factors, primarily the simple network architecture and the relatively small-scale dataset used for training. We chose a basic architecture, utilizing a ResNet-18 encoder and a straightforward decoder, to evaluate whether segmentation aids in AF identification. Despite the effectiveness of this approach, more sophisticated segmentation models could potentially yield significantly better results. Future work will employ advanced segmentation models, such as attention-based networks, to enhance segmentation outcomes. We will aim to develop a robust multi-task learning framework capable of simultaneously identifying AF and accurately segmenting infarct regions from brain MRI scans.

Addressing these challenges will not only enhance our model's segmentation capabilities but also contribute to advancing the field of automated diagnosis and treatment planning in post-stroke patients. Additionally, we aim to develop a multi-modal framework that integrates imaging and clinical information, typically used for AF risk assessment in patients, to create a robust AF risk assessment pipeline. While DWI and apparent diffusion coefficient (ADC) maps are commonly used together in clinical practices for stroke evaluation, this study relies solely on DWI images due to the unavailability of ADC maps. Future research can incorporate these additional imaging techniques to enhance the identification of AF-related strokes.

Our current research focused on LAA and AF, two prevalent stroke subtypes, while excluding other etiologies. These subtypes are clinically significant due to their higher recurrence rates and severity [5], [17]. Diagnosis of LAA typically involves vascular imaging, while identifying AF necessitates cardiac evaluation, contrasting with small artery occlusion (SAO), which can often be detected using brain MRI alone [4]. Moreover, targeted preventive measures for high-risk etiologies like LAA and AF generally result in more substantial risk reduction than lower-risk etiologies such as SAO [3], underscoring their importance for early intervention. Furthermore, the "other determined cause" category encompasses various complex stroke mechanisms that pose challenges for categorization and are therefore excluded from our study. We acknowledge that focusing on LAA and AF limits our study to specific stroke subtypes. However, this research represents the first attempt to establish a fully automated deep framework for the early identification of AF as the stroke etiology from brain MRI scans. By starting with these categories, which provide clear clinical and imaging markers, we ensure a solid foundation with reliable ground truth data. This approach lays the groundwork for future expansion to include other stroke mechanisms, thereby enhancing the framework's clinical relevance and utility.

VI. Conclusion

Atrial fibrillation (AF) is a prevalent cause of ischemic stroke, increasing recurrent stroke risk by two-fold. Early identification of AF during stroke onset is crucial for effective preventive strategies and treatment optimization. We introduced a multi-task learning framework using post-stroke DWI-MRI images for early detection of AF. Our deep learning approach efficiently identified AF, demonstrating a streamlined pipeline adaptable for analyzing brain MRI.

The proposed algorithm can seamlessly integrate into existing stroke diagnosis protocols as a risk stratification tool to identify individuals at risk of secondary strokes linked to AF. Importantly, our model aims to prioritize patients based on AF risk but does not replace current AF detection methods or serve as a definitive diagnostic tool. Patients identified as high-risk should undergo thorough cardiac monitoring for AF diagnosis. This facilitates targeted cardiac monitoring and definitive AF diagnosis where warranted and possible anticoagulation as clinically indicated. External validation is essential to assess our model's reliability and applicability across diverse datasets and clinical settings.

Acknowledgment

This research was supported by The University of Melbourne's Research Computing Services and the Petascale Campus Initiative.

References

[1] H. P. Adams Jr, B. H. Bendixen, L. J. Kappelle, J. Biller, B. B. Love, D. L. Gordon, and E. Marsh 3rd. Classification of subtype of acute ischemic stroke. definitions for use in a multicenter clinical trial. toast. trial of org 10172 in acute stroke treatment. *stroke*, 24(1):35–41, 1993. 2

[2] P. Amarenco, J. Bogousslavsky, L. Caplan, G. Donnan, and M. Hennerici. New approach to stroke subtyping: the asco (phenotypic) classification of stroke. *Cerebrovascular diseases*, 27(5):502–508, 2009. 2

[3] E. M. Arsava, J. Helenius, R. Avery, M. H. Sorgun, G.-M. Kim, O. M. Pontes-Neto, K. Y. Park, J. Rosand, M. Vangel, and H. Ay. Assessment of the predictive validity of etiologic stroke classification. *JAMA neurology*, 74(4):419–426, 2017. 3, 7

[4] H. Ay, T. Benner, E. Murat Arsava, K. L. Furie, A. B. Singhal, M. B. Jensen, C. Ayata, A. Towfighi, E. E. Smith, J. Y. Chong, et al. A computerized algorithm for etiologic classification of ischemic stroke: the causative classification of stroke system. *Stroke*, 38(11):2979–2984, 2007. 2, 7

[5] O. Y. Bang, P. H. Lee, S. Y. Joo, J. S. Lee, I. S. Joo, and K. Huh. Frequency and mechanisms of stroke recurrence after cryptogenic stroke. *Annals of neurology*, 54(2):227–234, 2003. 7

[6] J. P. Bretzman, A. S. Tseng, J. Graff-Radford, H.-C. Lee, S. J. Asirvatham, M. M. Mielke, D. S. Knopman, R. C. Petersen, C. R. Jack Jr, P. Vemuri, et al. Silent cerebral infarcts in patients with atrial fibrillation: Clinical implications of an imaging-adjusted cha2ds2-vasc score. *Cardiology Journal*, 29(5):766–772, 2022. 2

[7] I. R. Chavva, A. L. Crawford, M. H. Mazurek, M. M. Yuen, A. M. Prabhat, S. Payabvash, G. Sze, G. J. Falcone, C. C. Matouk, A. de Havenon, et al. Deep learning applications for acute stroke management. *Annals of Neurology*, 92(4):574–587, 2022. 2

[8] S. Chen, K. Ma, and Y. Zheng. Med3d: Transfer learning for 3d medical image analysis. *arXiv preprint arXiv:1904.00625*, 2019. 3

[9] D. Conen, N. Rodondi, A. Müller, J. H. Beer, P. Ammann, G. Moschovitis, A. Auricchio, D. Hayoz, R. Kobza, D. Shah, et al. Relationships of overt and silent brain lesions with cognitive function in patients with atrial fibrillation. *Journal of the American College of Cardiology*, 73(9):989–999, 2019. 2

[10] H. Cui, X. Wang, Y. Bian, S. Song, and D. D. Feng. Ischemic stroke clinical outcome prediction based on image signature selection from multimodality data. In *2018 40th Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC)*, pages 722–725, 2018. 2

[11] G. M. De Marchis, P. Krisai, L. Werlen, T. Sinnecker, S. Aeschbacher, T. D. Dittrich, A. A. Polymeris, M. Coslovksy, M. R. Blum, N. Rodondi, et al. Biomarker, imaging, and clinical factors associated with overt and covert stroke in patients with atrial fibrillation. *Stroke*, 54(10):2542–2551, 2023. 2

[12] L. Fast, U. Temuulen, K. Villringer, A. Kufner, H. F. Ali, E. Siebert, S. Huo, S. K. Piper, P. S. Sperber, T. Liman, et al. Machine learning-based prediction of clinical outcomes after first-ever ischemic stroke. *Frontiers in neurology*, 14:1114360, 2023. 2

[13] R. Feng, M. Badgeley, J. Mocco, and E. K. Oermann. Deep learning guided stroke management: a review of clinical applications. *Journal of neurointerventional surgery*, 10(4):358–362, 2018. 2

[14] L. Friberg, M. Rosenqvist, A. Lindgren, A. Terént, B. Norrving, and K. Asplund. High prevalence of atrial fibrillation among patients with ischemic stroke. *Stroke*, 45(9):2599–2605, 2014. 1

[15] K. G. Haeusler, D. Wilson, J. B. Fiebach, P. Kirchhof, and D. J. Werring. Brain mri to personalise atrial fibrillation therapy: current evidence and perspectives. *Heart*, 100(18):1408–1413, 2014. 2

[16] K. C. Ho, W. Speier, H. Zhang, F. Scalzo, S. El-Saden, and C. W. Arnold. A machine learning approach for classifying ischemic stroke onset time from imaging. *IEEE transactions on medical imaging*, 38(7):1666–1676, 2019. 2

[17] Y. Ko, S. Lee, J.-W. Chung, M.-K. Han, J.-M. Park, K. Kang, T. H. Park, S.-S. Park, Y.-J. Cho, K.-S. Hong, et al. Mri-based algorithm for acute ischemic stroke subtype classification. *Journal of stroke*, 16(3):161, 2014. 7

[18] C.-F. Liu, J. Hsu, X. Xu, S. Ramachandran, V. Wang, M. I. Miller, A. E. Hillis, and A. V. Faria. Deep learning-based detection and segmentation of diffusion abnormalities in acute ischemic stroke. *Communications Medicine*, 1(1):61, 2021. 2

[19] Y. Mayasi, J. Helenius, D. D. McManus, R. P. Goddeau, A. H. Jun-O'Connell, M. Moonis, and N. Henninger. Atrial fibrillation is associated with anterior predominant white matter lesions in patients presenting with embolic stroke. *Journal of Neurology, Neurosurgery & Psychiatry*, 89(1):6–13, 2018. 2

[20] A. Paszke, S. Gross, F. Massa, A. Lerer, J. Bradbury, G. Chanan, T. Killeen, Z. Lin, N. Gimelshein, L. Antiga, et al. Pytorch: An imperative style, high-performance deep learning library. *Advances in neural information processing systems*, 32, 2019. 5

[21] S. Penado, M. Cano, O. Acha, J. L. Hernández, and J. A. Riancho. Atrial fibrillation as a risk factor for stroke recurrence. *The American journal of medicine*, 114(3):206–210, 2003. 1

[22] J. R. Romero and P. A. Wolf. Epidemiology of stroke: legacy of the framingham heart study. *Global heart*, 8(1):67–75, 2013. 1

[23] R. M. Sarmento, F. F. X. Vasconcelos, P. P. R. Filho, W. Wu, and V. H. C. de Albuquerque. Automatic neuroimage processing and analysis in stroke—a systematic review. *IEEE Reviews in Biomedical Engineering*, 13:130–155, 2020. 2

[24] A. Sharobeam, L. Churilov, M. Parsons, G. A. Donnan, S. M. Davis, and B. Yan. Patterns of infarction on mri in patients with acute ischemic stroke and cardio-embolism: a systematic review and meta-analysis. *Frontiers in neurology*, 11:606521, 2020. 2

[25] S. A. Sheth, L. Giancardo, M. Colasurdo, V. M. Srinivasan, A. Niktabe, and P. Kan. Machine learning and acute stroke imaging. *Journal of neurointerventional surgery*, 15(2):195–199, 2023. 2

[26] R. Solovyev, A. A. Kalinin, and T. Gabruseva. 3d convolutional neural networks for stalled brain capillary detection. *Computers in Biology and Medicine*, 141:105089, 2022. 4

[27] Y. Yu, Y. Xie, T. Thamm, E. Gong, J. Ouyang, C. Huang, S. Christensen, M. P. Marks, M. G. Lansberg, G. W. Albers, et al. Use of deep learning to predict final ischemic stroke lesions from initial magnetic resonance imaging. *JAMA network open*, 3(3):e200772–e200772, 2020. 2

[28] P. A. Yushkevich, J. Piven, H. Cody Hazlett, R. Gimpel Smith, S. Ho, J. C. Gee, and G. Gerig. User-guided 3D active contour segmentation of anatomical structures: Significantly improved efficiency and reliability. *Neuroimage*, 31(3):1116–1128, 2006. 3

[29] R. Zhang, L. Zhao, W. Lou, J. M. Abrigo, V. C. T. Mok, W. C. W. Chu, D. Wang, and L. Shi. Automatic segmentation of acute ischemic stroke from dwi using 3-d fully convolutional densenets. *IEEE Transactions on Medical Imaging*, 37(9):2149–2160, 2018. 2

[30] Z. Zhang, K. Lin, J. Wang, L. Ding, Y. Sun, C. Fu, D. Qian, J. Li, and D. Huang. Searching for underlying atrial fibrillation using artificial intelligence-assisted mri images from ischemic stroke patients. *European Heart Journal*, 43(Supplement_2):ehac544–543, 2022. 2, 5

[31] N. Zungsontiporn and M. S. Link. Newer technologies for detection of atrial fibrillation. *Bmj*, 363, 2018. 1