

Received 7 November 2023, accepted 22 November 2023, date of publication 23 November 2023, date of current version 14 December 2023.

Digital Object Identifier 10.1109/ACCESS.2023.3336590

RESEARCH ARTICLE

CT Perfusion is All We Need: 4D CNN Segmentation of Penumbra and Core in Patients With Suspected Acute Ischemic Stroke

LUCA TOMASETTI¹, (Member, IEEE), KJERSTI ENGAN¹, (Senior Member, IEEE), LIV JORUNN HØLLESLI^{1,2}, KATHINKA DÆHLI KURZ^{1,2}, AND MAHDIEH KHANMOHAMMADI¹, (Member, IEEE)

¹Department of Electrical Engineering and Computer Science, University of Stavanger, 4021 Stavanger, Norway

²Stavanger Medical Imaging Laboratory (SMIL), Department of Radiology, Stavanger University Hospital, 4068 Stavanger, Norway

Corresponding author: Luca Tomasetti (luca.tomasetti@uis.no)

This work was supported by the Regional Ethic Committee under Project 2012/1499.

This work involved human subjects or animals in its research. Approval of all ethical and experimental procedures and protocols was granted by the Regional Ethic Committee.

ABSTRACT Stroke is the second leading cause of death worldwide, and around 87 % of strokes are ischemic strokes. Accurate and rapid prediction techniques for identifying ischemic regions, including dead tissue (core) and potentially salvageable tissue (penumbra), in patients with acute ischemic stroke (AIS) hold great clinical importance, as this can provide valuable information for diagnosis and treatment planning. Computed Tomography Perfusion (CTP) is often used as a primary tool for assessing stroke location, severity, and the volume of ischemic regions. Current automatic segmentation methods for CTP typically utilize pre-processed 3D parametric maps, traditionally used for clinical interpretation by radiologists. An alternative approach is to use the raw CTP data slice by slice as 2D+time input, where the spatial information over the volume is overlooked. Additionally, these methods primarily focus on segmenting core regions, yet predicting penumbra regions can be crucial for treatment planning. This paper investigates different methods to utilize the entire raw 4D CTP as input to fully exploit the spatio-temporal information, leading us to propose a 4D convolution layer in a 4D CNN network. Our comprehensive experiments on a local dataset of 152 patients divided into three groups show that our proposed models generate more precise results than other methods explored. Adopting the proposed *4D mJ-Net*, a Dice Coefficient of 0.53 and 0.23 is achieved for segmenting penumbra and core areas, respectively. Using the entire 4D CTP data for AIS segmentation offers improved precision and potentially better treatment planning in patients suspected of this condition.

INDEX TERMS 4D convolution, acute ischemic stroke, computed tomography perfusion, deep neural network, image segmentation.

I. INTRODUCTION

Neurological disorders are the primary contributor to disability-adjusted life years and the second most prevalent cause of death globally [1], with cerebral stroke as the leading cause of these disorders. This study focuses on ischemic stroke, which constitutes the majority of cerebral strokes [2].

The associate editor coordinating the review of this manuscript and approving it for publication was Hengyong Yu¹.

A precise and fast comprehension of the brain tissue affected by an ischemic stroke holds substantial value in guiding decision-making and treatment planning. An acute ischemic stroke (AIS) generally occurs if a segment of the supplying arteries of the brain is occluded by a blood clot and prevents the regular flow of oxygen-rich blood to the capillaries in the brain tissue.

The ischemic area can roughly be divided into two different types: 1) *penumbra*, areas where the tissue is still vital but

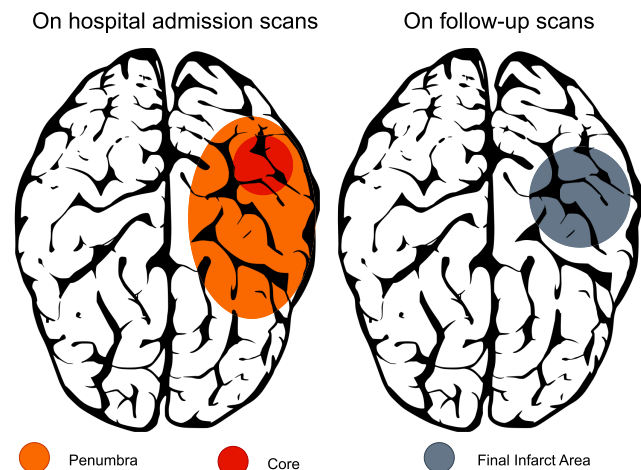


FIGURE 1. Visual representation of the various regions that can be defined after an IS: penumbra (in orange) and core (in red) can be localized during the initial stage of the neurological conditions, while the final infarct area (in dark gray) is determined during a follow-up study, after treatment.

critically hypoperfused [3]; and 2) *core*, referring to non-salvageable tissue. Fig. 1 represents a visual description of penumbra, in orange, and core, in red, after a scan at hospital admission. If blood flow is not restored timely, penumbra regions may develop rapidly into irreversibly damaged core regions. Therefore, a fast and accurate understanding of ischemic areas to plan the treatment and tailor further procedures to every single patient is fundamental.

The recommended modalities for diagnostic imaging in AIS patients are Computed Tomography (CT) and Magnetic Resonance Imaging (MRI) [4]. In the initial stages of an AIS, CT Perfusion (CTP) has proven to be a fast and beneficial tool for evaluating both diagnosis and prognosis [5]. MRI with Diffusion-weighted imaging (DWI) or non-contrast CT (NCCT) are commonly utilized after treatment to assess the *final infarct areas* (FIAs) [6] (in dark gray color in Fig. 1). These imaging modalities are obtained hours or days after the patient's treatment.

CTP is a four-dimensional (4D) spatio-temporal examination to assess the passage of blood in the brain. It is performed by acquiring a series of three-dimensional (3D) CT scans of a specific portion of the brain at time intervals during contrast agent injection. By using an iodinated contrast agent, density changes in the brain tissue over time can be analyzed. The shape and height of the time density curve depend on the brain tissue's perfusion [7]. The abundance of images in the raw 4D CTP poses a challenge for neuroradiologists in detecting ischemic strokes. The 4D volume of CTP requires simultaneous evaluation of the propagation of contrast fluid across all spatial regions of the specific portion of the brain over time to identify ischemic regions. However, this becomes impractical due to the complexity and the sheer volume of images involved. This complexity significantly prolongs the time required for analyzing the CTP study, which is highly detrimental in

situations requiring accurate diagnosis and prompt treatment decisions.

To overcome this challenge, medical doctors rely on software estimating a set of clinically interpretable parameters related to the propagation of the contrast fluid combining all the temporal information for each pixel generated from the 4D CTP scan. This gives a set of 3D parametric maps (PMs) [7], [8]. Commonly used PMs are cerebral blood flow (CBF), cerebral blood volume (CBV), time-to-maximum (TMAX), and time-to-peak (TTP). CBF represents the blood supply in the brain at a given time; CBV refers to the blood volume present at a given time in a brain region; TMAX is the flow-scaled residue function in the tissue, while TTP shows the time until the contrast agent reaches the tissue [7]. Maximum intensity projection (MIP) is also usually generated. MIP images are calculated as the maximum Hounsfield unit (HU) value over the time sequence of the CTP, providing a 3D volume from the 4D acquisition of CTP. Although PMs provide helpful information about ischemic brain tissue, extracting them from the 4D CTP scans limits the spatio-temporal information only to specific subsets of information [9].

In this study, the objective is to develop a fully automated method to segment the penumbra and core regions in AIS patients based on the raw 4D CTP input. The main contributions of this work can be summarized in three points:

- 1) We propose a 4D convolution layer and use that to propose a DNN model, *4D mJ-Net*, to segment ischemic core and penumbra areas from 4D CTP scans.
- 2) We extend multiple existing methods for 3D CTP to 4D CTP to perform a comparison with the proposed 4D convolution solution.
- 3) To assess the results, we use manual annotations obtained by two expert neuroradiologists from the 4D CTP data upon patients' admission. We also demonstrate the feasibility of our proposed methods by comparing their performances with existing models that rely on different inputs.

II. PREVIOUS WORK

Several methods [6], [10], [11], [12], [13] have used thresholding techniques to predict the ischemic areas from the PMs. However, simple thresholding approaches over-simplify the complexity in AIS [14], [15].

In the past years, Deep Neural Networks (DNNs), and especially Convolutional Neural Networks (CNNs), have been successfully applied in numerous medical applications: image classification tasks [16], [17], [18], [19], [20], automatic video analysis [21], [22], and activity recognition [23], [24], [25]. Automatic image segmentation adopting U-Net structure [26] and its numerous variants have produced innovative outcomes for several applications [27], [28], [29], [30], [31].

Several DNNs have been proposed for AIS applications to predict and segment *only* the FIAs using CT studies in combination with PMs derived from CTP scans as input

[32], [33], [34], [35], [36]. Other researchers have proposed architectures to segment the ischemic lesion (i.e., the core) from the images obtained at hospital admission. Kasasbeh et al. [37] were the first to implement a CNN with a set of PMs as input for ischemic core segmentation. Tomasetti et al. [38] proposed a few-shot self-supervised architecture for hypoperfused (core + penumbra) tissue segmentation using a combination of PMs and raw scans as input of the model. The work demonstrated the feasibility of using self-supervised techniques to segment this tissue type. Werdiger et al. [39] introduced a machine learning segmentation method to delineate hypoperfused tissue, demonstrating the capabilities of this methodology over classic thresholding approaches. They used four PMs as input features for their model. However, a general problem with all the methods mentioned above is relying on commercial CTP software and using heavily pre-processed information (i.e., PMs) rather than taking advantage of the totality of the raw 4D CTP scans.

DNNs are more suitable for discovering information from raw data. Nevertheless, relying on raw data (directly exploiting the temporal and spatial dimensions) is scarcely explored in the literature for AIS applications. The task is challenging because of the low contrast and low signal-to-noise ratio in the CTP scans. Relatively few studies proposed DNN models with encouraging results, exploiting the temporal dimension to assess acute stroke lesions using 4D CTP scans. Soltanpour et al. [9] utilized CTP images to create 2D matrices in which each row is a voxel, and each column is a time point. The 2D matrices are used as input for a model that shows encouraging results in differentiating healthy tissue from FIAs. Vries et al. [40] promoted a 2D+time symmetry-aware CNN-based architecture to segment FIAs using solely CTP scans. Their work estimated the irreversibly damaged areas, demonstrating the possibilities of using 4D CTP images for this task. Bertels et al. [41] used a U-Net-like structure for segmenting FIAs using CTP scans as an input plus contra-lateral information. Results were promising, but further research is needed due to their far-from-ideal registration of the contra-lateral information. Rosa et al. [42] introduced a two-step model for estimating FIAs using the 4D CTP series as input. They first generate an arterial input function and later deconvolve it with a singular value decomposition approach to find the infarction. Amador et al. [43] designed a framework based on the Temporal Convolution Network to predict AIS FIAs from 4D CTP studies. Due to memory constraints, they independently processed each 2D slice of the 4D CTP dataset. In their recent work, Amador et al. [44] also proposed an extension of their model where 3D+time tensors of the ipsilateral stroke hemisphere are used as input to predict FIAs. Robben et al. [45] proposed a DNN that predicts the FIAs directly from 4D raw CTP plus patients metadata. Their proposed architecture relied on a series of 3D Convolution layers; the input is a list of 4D CTP scans sampled at different resolutions. Their method presented promising segmentation results; however, their main target was to estimate the final infarct volume, allowing clinicians

to simulate different treatments and gain insight into the procedures. They were not taking into consideration the penumbra in their study. Plus, the quality of the ground truth images is debatable since they rely on NCCT follow-up images acquired between 24 hours and five days after patient's admission. It has been reported that FIAs can grow after 24 hours in NCCT measurements [46].

All the segmentation methods in the previous paragraph rely on ground truth labels obtained from DWI and/or NCCT hours or days after the patient's admission since they predict FIAs. Even though follow-up images (DWI and NCCT) represent the gold standard for estimating core [6], there are some limitations with these techniques [47], [48]. Follow-up images can only be used to assess FIAs but not penumbra regions. Plus, some studies have demonstrated that the detected FIAs can be partially reverse in DWI performed in an early time window [48], [49], [50].

Previous studies have indicated the potential of 4D data in AIS prediction [9], [40], [41], [42], [43], [44], [45], with a critical gap as they only consider predicting FIAs. An appropriate method is still required to simultaneously handle the spatio-temporal information for segmenting the ischemic *core* and *penumbra* regions. Understanding the penumbra's extension during the ischemic stroke's first stages is crucial for treatment decision [51], [52]. To the best of our knowledge, our work in Tomasetti et al. [52] using machine learning, and later in Tomasetti et al. [53], [54] using DNN, were the first and only to segment *both* core and penumbra areas. In [52] and [54], the PMs were used as input, and in [53] 2D + time CTP images were segmented slice-by-slice.

Building upon our previous works [52], [53], [54] and filling the critical gap left by previous studies [9], [40], [41], [42], [43], [44], [45], in this paper, we present and investigate three novel models to segment the two ischemic regions (core and penumbra), where the input is the entire 4D CTP scans arranged in different ways to exploit the spatio-temporal nature of the data. We compare all models with previous work based on PMs [54] and slice-by-slice CTP [53], and two methods from the literature proposed by Amador et al. [43], [44].

III. DATA MATERIAL

A section of the brain is repeatedly scanned during the passage of 40 ml iodine-containing contrast agent (Omnipaque 350 mg/ml) and 40 ml isotonic saline in a cubital vein with a flow rate of 6 ml/s to highlights changes in the tissue; the scan delay was four seconds. Each brain slice contains a fixed number of time points t_{\max} representing the temporal dimension. The width and height of each image are 512×512 pixels with a resolution of 0.4258 mm/pixel and a slice thickness of 5mm. The first twenty time points are acquired every 1s, and the remaining ten images every 2s.

CTP scans from 152 patients collected between January 2014 and August 2020 formed the dataset. 137 of these patients had an AIS with a visible perfusion deficit. During the diagnostic workup, the remaining 15 patients were

admitted with suspected strokes but were determined not to have suffered from a stroke episode after the diagnostic workup. The raw CTP studies are saved as DICOM files. Raw perfusion data from the CTP examination was used to generate PMs with the software “syngo.via” from Siemens Healthineers, with manufacturer default settings. The arterial input function was automatically selected, with few exceptions where it was chosen manually (e.g., severe cardiac failure).

The patients were divided into the following groups: 77 patients with large vessel occlusion (LVO), 60 patients with non-large vessel occlusion (Non-LVO), and the remaining 15 patients without ischemic stroke (WIS). Based on CT angiography, LVO was defined as occlusion of any of the following arteries: the internal carotid artery, M1 and proximal M2 segment of the middle cerebral artery, A1 segment of the anterior cerebral artery, P1 segment of the posterior cerebral artery, basilar artery, and vertebral artery occlusion. Non-LVO was defined as patients with perfusion deficit with more distal artery occlusion or with perfusion deficit without visible artery occlusion. The dataset is randomly split into a training, validation, and test set. The percentage of the three subsets (LVO, Non-LVO, WIS) is equally distributed among the sets, as shown in Table 1.

TABLE 1. Division in training, validation, and test dataset.

	Training (#; %)	Validation (#; %)	Test (#; %)	Tot. (#; %)
LVO	42; 54.5	16; 20.8	19; 24.7	77; 50.6
Non-LVO	36; 60	13; 21.7	11; 18.3	60; 30.5
WIS	9; 60	3; 20	3; 20	15; 9.8
Total	87; 58.6	32; 19.7	33; 21.7	152; 100

A. GROUND TRUTH

The manual annotations are based on the entire CT dataset, including the PMs derived from CTP. MRI performed during the first days after hospital admission was also utilized. Two expert neuroradiologists manually annotated ground truth images by utilizing the complete set of the CT examination (NCCT, CT angiography, and CTP), which includes PMs from the CTP (CBV, CBF, TTP, TMAX) and the MIP images. The PMs were visually assessed. In general, ischemic regions with increased TTP and TMAX and reduced CBF but preserved CBV were considered as penumbra, while areas with additionally reduced CBV were deemed as core. Additionally, the MRI examination, including DWI, obtained within 1 to 3 days after the CT examination, and clinical information, was used to assist in generating the ground truth images. The annotations were performed using an in-house developed software in Matlab.¹

¹The code is publicly available at the following link <https://github.com/Biomedical-Data-Analysis-Laboratory/CTP-Matlab>

IV. BACKGROUND THEORY

A. NOTATION

Table 2 presents the various formal notations adopted in the remainder of the paper. Let the data obtained from a CTP scan be defined as a 4D tensor $V \in \mathbb{R}^{(X \times Y \times Z \times T)}$. After a series of pre-processing steps (details in Sec. IV-B), we define the 4D tensor as $\tilde{V} \in \mathbb{R}^{(X \times Y \times Z \times T)}$. The four dimensions of a CTP scan are defined as width (X), height (Y), depth (Z), and time (T). The list of time points in the time dimension is given by $t = [t_j | \forall j \in \{1, 2, \dots, t_{\max}\}]$, where t_{\max} is the last time point of the list. We indicate how the notation superscript adopts the time dimension in the various inputs. Furthermore, we define $z = [z_i | \forall i \in \{1, 2, \dots, z_{\max}\}]$ as the list of brain slices in the depth dimension, where z_{\max} corresponds to the last slice. We illustrate how the depth dimension is being used in the inputs through the notation subscript. Fig. 3 displays the input combination of all the techniques.

TABLE 2. List of formal notations used in the paper.

Notation	Description
$t = [t_j \forall j \in \{1, 2, \dots, t_{\max}\}]$	List of all the time points.
t_{\max}	Last time point in the time dimension.
$z = [z_i \forall i \in \{1, 2, \dots, z_{\max}\}]$	List of all the slices.
z_{\max}	Last slice in the depth dimension.
$\mathcal{I} = \{i - 1, i, i + 1\}$	Set of indexes i , plus its neighbours $i - 1$ and $i + 1$.
$z_{\mathcal{I}} = \{z_{i-1}, z_i, z_{i+1}\}$	Set of slices: z_{i-1}, z_i, z_{i+1} .
$V \in \mathbb{R}^{(X \times Y \times Z \times T)}$	4D raw CTP input.
$\mathcal{C} = \{\text{healthy brain, penumbra, core}\}$	Set of classes.
$I_{z_i}^t \in \mathbb{R}^{(X \times Y)}$	2D brain slice z_i at time point t_j .
$P_{z_i} \in \mathbb{R}^{(X \times Y)}$	2D probability output of brain slice z_i .
\sim	Input after pre-processing steps.
$\bar{\cdot}$	List of inputs
$\varphi(\cdot)$	Concatenation function.
$\hat{\cdot}$	Concatenated inputs after passing through $\varphi(\cdot)$.
$\tilde{V}_{z_i}^t = [\tilde{V}_{z_i}^t \forall t_j \in t] \in \mathbb{R}^{(X \times Y)}$	List of 2D images $I_{z_i}^t$ for all the time points t . Input for 2D-TCN (Sec. V-A3).
$\hat{V}_{z_i}^t = \varphi(\tilde{V}_{z_i}^t \forall t_j \in t) \in \mathbb{R}^{(X \times Y \times T)}$	2D+time volume of slice z_i . Input for mJ-Net (Sec. V-A2).
$\hat{V}_{z_{\mathcal{I}}}^t = \varphi(\hat{V}_{z_{i-1}}^t, \hat{V}_{z_i}^t, \hat{V}_{z_{i+1}}^t) \in \mathbb{R}^{(X \times Y \times Z)}$	3D volume of slices $z_{\mathcal{I}}$ for a time point t_j .
$\tilde{V}_{z_{\mathcal{I}}}^t = [\tilde{V}_{z_{\mathcal{I}}}^t \forall t_j \in t] \in \mathbb{R}^{(X \times Y \times Z)}$	List of 3D volumes of slices $z_{\mathcal{I}}$ for all the time points t . Input for 3D-TCN (Sec. V-B1).
$\hat{V}_{z_{\mathcal{I}}}^t = [\hat{V}_{z_{i-1}}^t, \hat{V}_{z_i}^t, \hat{V}_{z_{i+1}}^t] \in \mathbb{R}^{(X \times Y \times Z \times T)}$	List of 2D+time volumes for slices $z_{\mathcal{I}}$. Input for 3D+time mJ-Net (Sec. V-B2).
$\hat{V}_{z_{\mathcal{I}}}^t = \varphi(\hat{V}_{z_{i-1}}^t, \hat{V}_{z_i}^t, \hat{V}_{z_{i+1}}^t) \in \mathbb{R}^{(X \times Y \times Z \times T)}$	4D Tensor of slices $z_{\mathcal{I}}$ over all the time points t . Input for 4D mJ-Net (Sec. V-B3).

All methods return a 3D output, segmenting the images P_{z_i} slice-by-slice. The segmented 2D image P_{z_i} corresponds to a brain slice z_i at index i . The predicted image P_{z_i} contains brain tissue segmented with the classes \mathcal{C} (if any): healthy brain, penumbra, and core.

B. PRE-PROCESSING STEPS

The 4D CTP dataset underwent a series of pre-processing steps to extract brain tissue from the raw CTP scans.

Algorithm 1 Pre-Processing Steps for One Single Patient Study

Input: 4D CTP scan: $V(x, y, z, t)$
 ref $\leftarrow V(x, y, z, t_1)_{\text{HU}}$ {Get the 1st time point image as frame of reference}
for $j = 2$ to t_{max} **do**
 Co-register $V(x, y, z, t_j)$ using ref as frame of reference.
end for
 $V(x, y, z, t)_{\text{HU}} \leftarrow \text{ConvertToHU}(V(x, y, z, t))$
 $\hat{V}(x, y, z, t) \leftarrow \text{BrainExtraction}(V(x, y, z, t)_{\text{HU}})$ {The brain extraction function is designed by Najm et al. [55]}
 $\check{V}(x, y, z, t) \leftarrow \text{GammaCorrection}(\hat{V}(x, y, z, t))$
 $z_{\text{high}} \leftarrow \text{GetSliceWithHighestIntensityValue}(\check{V}(x, y, z, t))$
 $\text{bins} \leftarrow 2^{16}$
 $\tilde{V}(x, y, z_{\text{high}}, t), \mathcal{T}_{z_{\text{high}}} \leftarrow \text{HistEq}(\check{V}(x, y, z_{\text{high}}, t), \text{bins})$
 { $\mathcal{T}_{z_{\text{high}}}$ is the grayscale transformation for z_{high} }
 $\tilde{V}(x, y, z_{\text{high}}, t) \leftarrow \frac{\check{V}(x, y, z_{\text{high}}, t) - \text{mean}(\check{V}(x, y, z_{\text{high}}, t))}{\sigma(\check{V}(x, y, z_{\text{high}}, t))}$ {Standardization of the data}
for $i = 1$ to z_{max} **do**
 if $i \neq z_{\text{high}}$ **then**
 $\check{V}(x, y, z_i, t) \leftarrow \text{HistEq}(\check{V}(x, y, z_i, t), \mathcal{T}_{z_{\text{high}}})$
 $\tilde{V}(x, y, z_i, t) \leftarrow \frac{\check{V}(x, y, z_i, t) - \text{mean}(\check{V}(x, y, z_i, t))}{\sigma(\check{V}(x, y, z_i, t))}$
 end if
end for
return Processed 4D CTP scan: $\tilde{V}(x, y, z, t)$

Algorithm 1 describes in detail the various steps for pre-processing all the data in a patient study. Furthermore, each step can be summarized as follows:

- 1) Co-registration of all the images in the 4D CTP scan using the first time point image as the frame of reference in order to correct possible motion artifacts. An intensity-based image registration with similarity transformation was used in this step.
- 2) All the registered CTP scans were encoded into HU values to have a known quantitative scale to describe radiodensity efficiently. We used the following equation to calculate the HU value for a voxel V with a rescale slope (RS) and a rescale intercept (RI) extracted from the DICOM header: $V(x, y, z, t)_{\text{HU}} = V(x, y, z, t) \cdot \text{RS} + \text{RI}$
- 3) Brain extraction of CT studies plays an essential role in stroke imaging research [55], [56]. An automatic brain extraction method designed by Najm et al. [55] was selected for this purpose due to its proven efficiency with CT datasets and public availability.
- 4) Gamma correction ($\gamma = 0.5$) and histogram equalization (HE) were also performed after step 3 to increase contrast and visual interpretability on the CTP scans.
- 5) Finally, z-score (z) on the enhanced 4D tensor is applied to normally distribute the data. Thus, every study has a mean of 0 and a standard deviation of 1, ensuring that all CTP studies have a similar scale.

The input for all the methods (except for the *Multi-input PMs's* approach) follows the same pre-processing steps. These steps were performed to improve the quality of the images by enhancing the contrast. For a detailed explanation of the pre-processing steps, we refer the reader to [38] and [53]. An additional re-sampling step for all the images was performed to ensure uniform distribution in the temporal dimension. The pre-processing steps and re-sampling effects are examined in an ablation study in Sec. VI-C.

C. CONVOLUTION IN MANY DIMENSIONS

A handful of DNN methods have been proposed to exploit 4D data with a full 4D Convolution (4D-Conv) layer. Using 4D-Conv for 4D data produced better performances than using multiple 3D Convolutions (3D-Conv) on the same data. Gessert et al. [57] and Bengs et al. [58] proposed a 4D spatio-temporal convolutional network to optical coherence tomography force estimation. They demonstrated that using the full 4D data information yields better performances than 3D data. Myronenko et al. [59] introduced a 4D CNN to segment cardiac volumetric sequences using CT scans, showing the advantages of using their proposed architecture compared to a classic 3D CNN.

In the remainder of this manuscript, let's define $I(x, y, z, t) \in \mathbb{R}^4$, $\mathcal{H}(w, h, d, p) \in \mathbb{R}^4$ as a 4D tensor and a 4D kernel, respectively. The x and w indicate the width of the 4D structures; y and h express the height dimension; z and d define the depth dimension, while t and p represent the time dimension of the 4D structures. Like a 3D-Conv can be represented as the sum of multiple 2D Convolution (2D-Conv) along the depth dimension, a 4D-Conv operation can be described as the sum of numerous 3D-Conv along the temporal dimension. The loop rearrangement to avoid repeated 3D-Conv operations allows a true (non-separable) 4D convolution operation [59].

A 4D-Conv $g''''(x, y, z, t)$ of a 4D input $I(x, y, z, t) \in \mathbb{R}^4$ and a 4D kernel $\mathcal{H}(w, h, d, p) \in \mathbb{R}^4$ can be defined as:

$$g''''(x, y, z, t) = \mathcal{H}(w, h, d, p) \otimes I(x, y, z, t)$$

$$= \sum_{l=0}^{p-1} \sum_{k=0}^{d-1} \sum_{i=0}^{w-1} \sum_{j=0}^{h-1} \mathcal{H}(i, j, k, l) I(\tilde{x}, \tilde{y}, \tilde{z}, \tilde{t}),$$

where $\tilde{x} \equiv x + \tilde{w} - i$, $\tilde{y} \equiv y + \tilde{h} - j$, $\tilde{z} \equiv z + \tilde{d} - k$, and $\tilde{t} \equiv t + \tilde{p} - l$. For a detailed description of the 2D-Conv and 3D-Conv, we refer to the additional supplemental material linked with the manuscript.

We define a 4D-Conv $g''''(x, y, z, t)$ and a 4D kernel $\mathcal{H}(w, h, d, p)$ as the sum of multiple 3D-Conv over a specific dimension, i.e., the third dimension. The 4D kernel $\mathcal{H}(w, h, d, p)$ can be seen as a list of $\mathcal{H}(w, h, k, p) \forall k \in d-1$, where $\mathcal{H}(w, h, k, p)$ is a 2D+time volume of the k th slice over the entire p elements in the temporal dimension:

$$g''''(x, y, z, t) = \mathcal{H}(w, h, d, p) \otimes I(x, y, z, t)$$

$$= \sum_{k=0}^{d-1} \mathcal{H}(w, h, k, p) I(x, y, z + \tilde{d} - k, t)$$

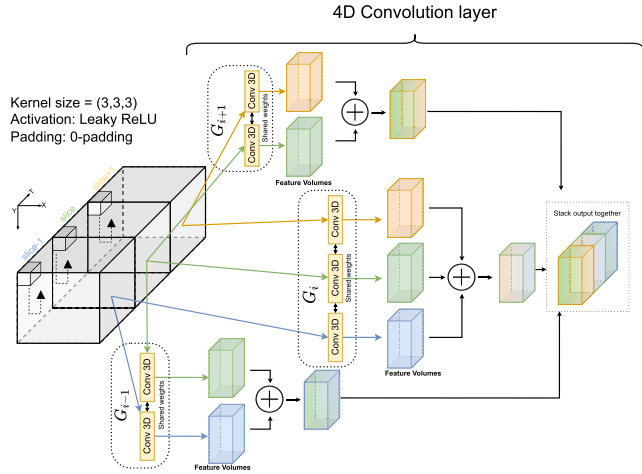


FIGURE 2. Visual representation of a 4D-Conv layer. The input is $\hat{V}_{z_{\mathcal{I}}}^t = \varphi(\hat{V}_{z_{i-1}}^t, \hat{V}_{z_i}^t, \hat{V}_{z_{i+1}}^t) \in \mathbb{R}^{(X \times Y \times Z \times T)}$. A series of 3D-Conv operations are calculated over a 4D input. Several groups (G_{i-1} , G_i , G_{i+1}) are generated, one for each volume involved in the operation.

$$= \sum_{k=0}^{d-1} g'''(x, y, z + \tilde{d} - k, t),$$

where $I(x, y, z + \tilde{d} - k, t) \in \mathbb{R}^3$ is a 2D+time volume at slice $z + \tilde{d} - k$ over the totality t element in the temporal dimension, and $[z + \tilde{d} - k \in z_m | \forall m \in \{1, \dots, z_{\max}\}]$

Fig. 2 visually explains the proposed 4D-Conv layer. The input for our 4D-Conv layer is a 4D tensor $\hat{V}_{z_{\mathcal{I}}}^t$, where $z_{\mathcal{I}}$ is a set of neighboring slices $\{z_{i-1}, z_i, z_{i+1}\}$ (Table 2). The 2D+time volume of the i th brain slice $\hat{V}_{z_i}^t$ over all the time points t , and the two 2D+time volumes of the neighboring brain slices ($\hat{V}_{z_{i-1}}^t$ and $\hat{V}_{z_{i+1}}^t$). The 4D-Conv layer uses a loop rearrangement with three 3D-Conv groups (G_{i-1} , G_i , G_{i+1}), one for each volume slice involved in the operation. In each group G_i , several 3D-Conv layers are created. All convolution layers in each group shared the weights. Each 3D-Conv layer is used for a single input volume. The number of layers depends on the legal subscript indexes, i.e., for the group G_{i-1} , there are two 3D-Conv layers since the legal subscript indexes are $\{i-1, i\}$; the indexes are given by the current volume slice z_{i-1} and the only neighboring volume slice z_i . The output of each group is a set of feature volumes summed together. The resulting feature volumes are stacked together to keep the same dimension as the input.

V. EXISTING METHODS AND PROPOSED 4D METHODS

This paper presents three novel deep learning (DL) approaches that accommodate 4D input data. In the remainder of the paper, they are called *3D-TCN* (Sec. V-B1), *3D+time mJ-Net* (Sec. V-B2), and *4D mJ-Net* (Sec. V-B3).

Together with the proposed approaches, we implemented and compared the *2D-TCN* [43] (Sec. V-A3), the *3D-TCN-SE* [44] (Sec. V-A4), and the *mJ-Net* [53] (Sec. V-A2) to validate the performances of our models. We also compare the models with a method that uses a set of PMs as input [54]. In the

remainder of the paper, we call this architecture *Multi-input PMs* (Sec. V-A1). Fig. 3 visually compares the input utilized for the various approaches.

A. EXISTING METHODS

1) APPROACH 1: MULTI-INPUT PMS

The *Multi-input PMs* model was proposed in [54]. This architecture was used as a baseline study because all the PMs were input for the model. The input for the architecture is a list of PMs for each brain slice z_i : PMs_{z_i} , as shown in Table 3. The loss function implemented is the Focal Tversky loss (FTL) [60]; for a specific class c , the FTL is defined as:

$$FTL(x, y) = \sum_c^C (1 - TI_c)^{1/\gamma},$$

where $\gamma \geq 1$ is a hyper-parameter that forces the loss function to focus more on less accurate predictions that have been misclassified [60]. Denoting $x_{i,c} \in [0, 1]$ as the probability of the i th predicted pixel to belong to class c ; $y_{i,c} \in \{0, 1\}$ as the pixel i with class c in a ground truth image, TI_c is the Tversky index (TI) for a class c defined as:

$$TI_c = \frac{\sum_{i=1}^{M \times N} x_{i,c} y_{i,c}}{\sum_{i=1}^{M \times N} x_{i,c} y_{i,c} + \alpha \sum_{i=1}^{M \times N} \widehat{x}_{i,c} y_{i,c} + \beta \sum_{i=1}^{M \times N} x_{i,c} \widehat{y}_{i,c}},$$

where $\widehat{x}_{i,c} = 1 - x_{i,c}$ is the probability that the i th pixel is not of class c , and $\widehat{y}_{i,c} = 1 - y_{i,c}$ represents the complement of pixel i in a ground truth image. The hyper-parameters α and β control the trade-off between precision and recall. We refer the reader to [54] for a more extensive explanation and discussion about this approach.

2) APPROACH 2: MJ-NET

The *mJ-Net* approach was proposed in [53]. As presented in Table 3, the input $\hat{V}_{z_i}^t(x, y)$ for *mJ-Net* is a 2D+time volume of the same brain slice z_i at index i over all the time points t . We define the dimension of this input as 2D+time; the first dimension of the input is time.

The loss function used for this method is the soft Dice Coefficient loss (SDCL) [61]. The SDCL is a modified version of the Dice Coefficient score mainly used in medical domains where the classes to predict are highly unbalanced due to a small region of interest compared to the background of the scans. The SDCL can be written as:

$$SDCL(x, y) = \sum_c^C \left(1 - \frac{2 \sum_i^{M \times N} x_{i,c} y_{i,c}}{\sum_i^{M \times N} x_{i,c}^2 + \sum_i^{M \times N} y_{i,c}^2} \right)$$

The first section of the model contains 3D-Conv layers to extract information from the temporal dimension, while the second part follows the classic U-Net structure [26]. For more details about the *mJ-Net* approach, we refer the reader to [53].

3) APPROACH 3: 2D TEMPORAL CONVOLUTIONAL NETWORK

For comparison reasons, we implemented the method proposed by Amador et al. [43]. We call this architecture

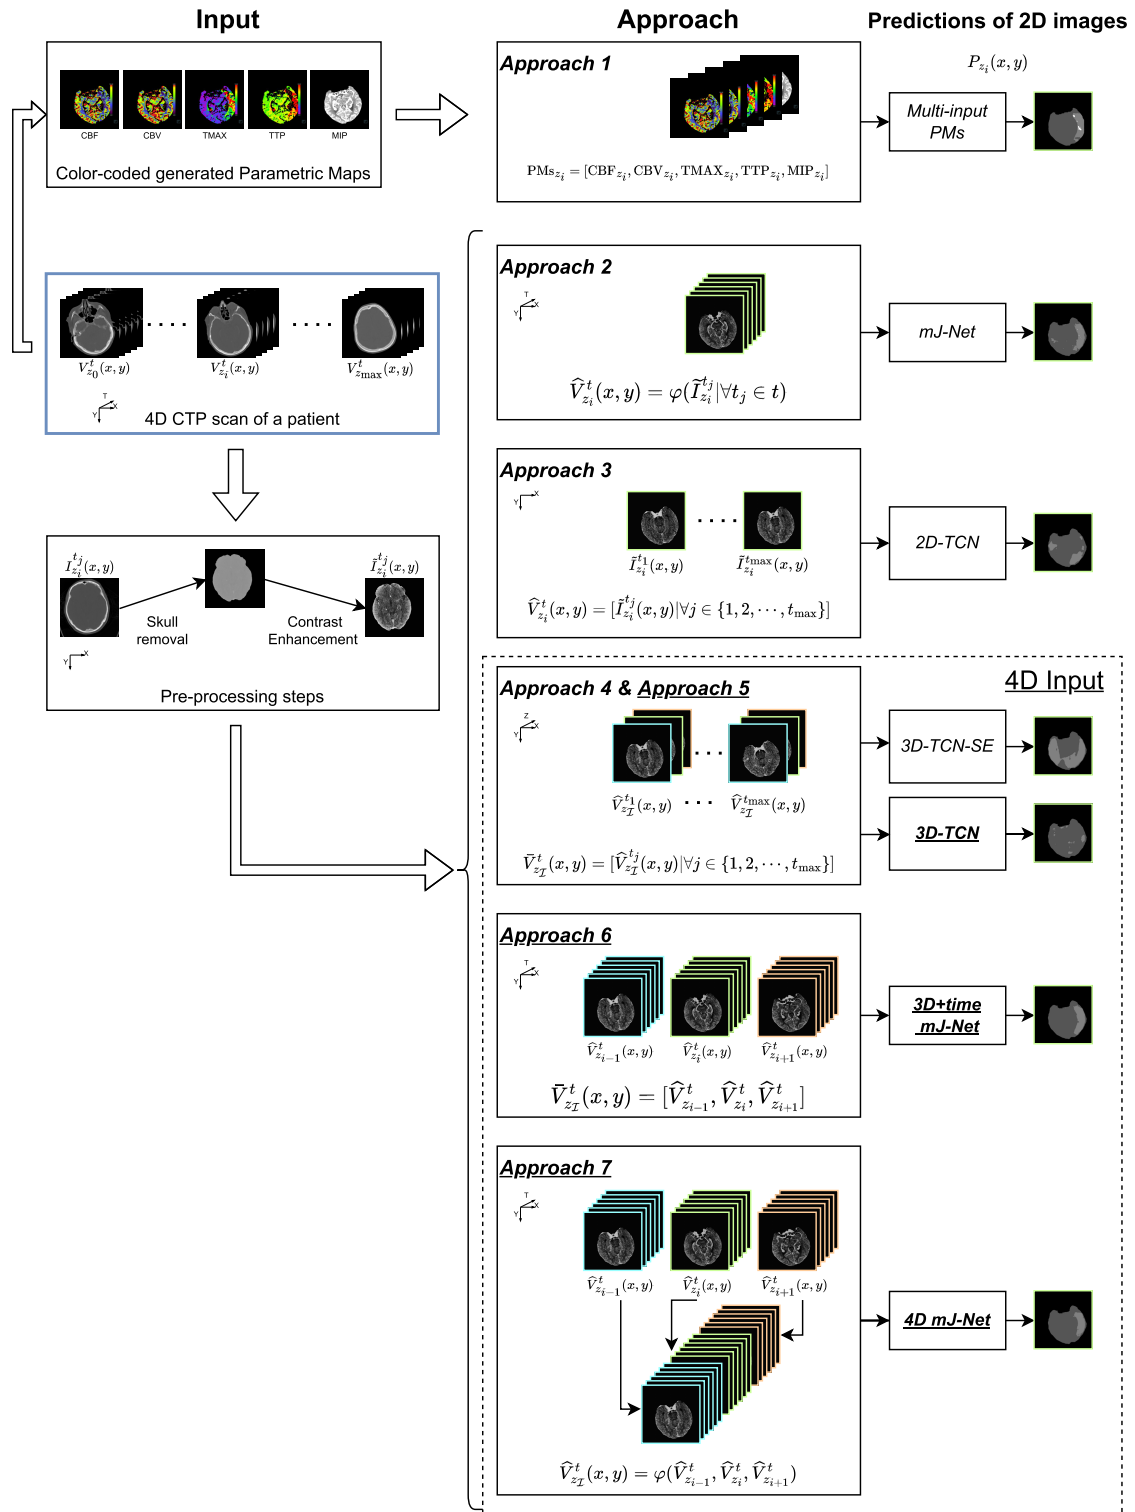

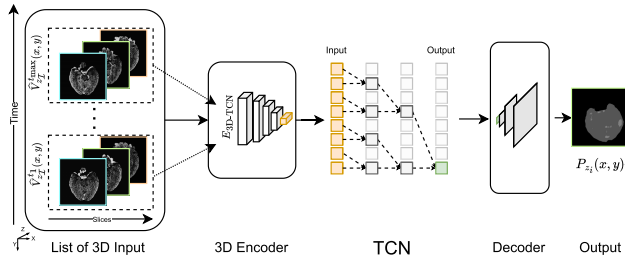


FIGURE 3. Visual comparison of the input for each implemented approach. Every 4D CTP patient’s study $V \in \mathbb{R}^{(X \times Y \times Z \times T)}$ undergoes a series of pre-processing steps to enhance each CTP scan. Approach 1 (*Multi-input PMs*) [54] accepts a list of PMs generated from a CTP study in input. Approach 2 (*mJ-Net*) [53] use a 2D+time volume $\hat{V}_{z_i}^t(x, y)$ as input. Approach 3 (*2D-TCN*) and Approach 4 (*3D-TCN-SE*) follows the model proposed by Amador et al. [43] and Amador et al. [44], respectively. The resulting 4D tensor is fed to one of the approaches. The proposed approach 5 (*3D-TCN*), approach 6 (*3D+time mJ-Net*), and approach 7 (*4D mJ-Net*) take in input the entire 4D CTP processed data \tilde{V} .

2D-TCN in the remainder of the paper. Fig. 4a provides a general architecture overview. They proposed a Temporal Convolutional Network (TCN), which has been shown to outperform conventional neural networks in different tasks [62]. Moreover, a TCN has a lower memory requirement for training than other Recurrent Neural Networks [62].



(a) 2D-TCN architecture proposed by Amador *et al.* [43].




(b) 3D-TCN-SE proposed by Amador *et al.* [44]

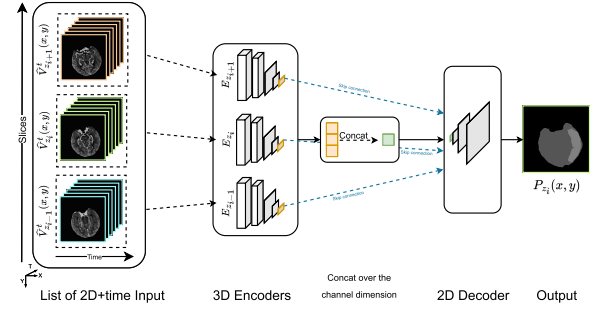
FIGURE 4. Visual comparison between (a) the 2D-TCN architecture [43] and (b) the 3D-TCN-SE [44].

The 2D-TCN was trained with the exact implementation as the original work (see Table 3 for more details). The 2D-TCN model receives the 4D CTP scans in input re-sampled to 1 second per time point. The 4D input is processed as a list of 2D brain slices z_i for each time point t . Thus, the actual input for the 2D-TCN is a list $\bar{V}_{z_i}^t$, as mentioned in Table 3. The list $\bar{V}_{z_i}^t$ contains all the time points of the brain slice z_i . Every 2D input image of the list $\bar{I}_{z_i}^{t_j}$ at time point t_j is fed to a 2D encoder $E_{2D-TCN}^{t_j}$ to extract features in the latent space. Each $E_{2D-TCN}^{t_j}$ encoder returns a $(4 \times 4 \times \text{Ch})$ feature vector, where Ch corresponds to the number of channels. The architecture merges the low-level feature vectors across the different t_j time points to capture the spatio-temporal information. The merged feature vector $\text{ETOT}_{2D-TCN} = [E_{2D-TCN}^{t_j} | \forall t_j \in t]$ is used as input to the TCN, which yields a one-dimensional vector O_{2D-TCN} of 64 elements. Finally, a decoder takes the O_{2D-TCN} and generates a final 2D image $P_{z_i}(x, y)$. The Dice Coefficient loss (DCL), the same as the original paper, was implemented as the loss function as follows:

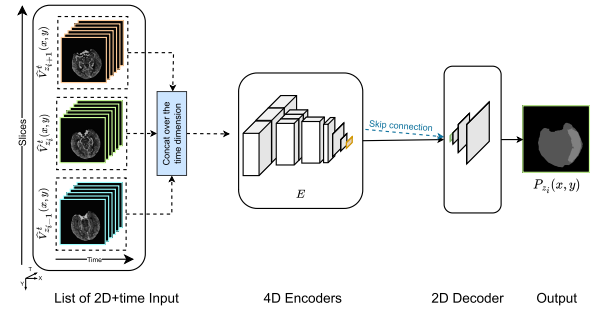
$$\text{DCL}(x, y) = \sum_c \left(1 - \frac{2 \sum_i^{M \times N} x_{i,c} y_{i,c}}{\sum_i^{M \times N} x_{i,c} + \sum_i^{M \times N} y_{i,c}} \right)$$



(a) 3D-TCN architecture (Sec. V-B1)



(b) 3D+time mJ-Net architecture (Sec. V-B2)



(c) 4D mJ-Net architecture (Sec. V-B3)

FIGURE 5. Visual comparison between the proposed architectures: (a) the 3D-TCN, (b) the 3D+time mJ-Net, and (c) 4D mJ-Net.

For more details about the 2D-TCN approach, we refer the reader to [43].

4) APPROACH 4: 3D TEMPORAL CONVOLUTIONAL NETWORK SINGLE ENCODER

We implemented a similar method from Amador *et al.* [44]. We call this approach 3D-TCN-SE due to using a single encoder (SE). Fig. 4b shows a simplified version of the proposed architecture, emphasizing the input difference between the 2D-TCN and this model.

The 3D-TCN-SE model receives the 4D CTP scans re-sampled to 1 second per time point. The input is a list of t 3D volumes $\bar{V}_{z_i}^t$ (Table 3). Each 3D input volume in the list $\bar{V}_{z_i}^{t_j}(x, y)$ corresponds to the concatenation of the i th brain slice z_i plus its neighbouring slices z_{i-1} and z_{i+1} over a specific time point t_j . The 3D-TCN-SE approach uses a single encoder E_{3D-TCN} for all the elements in the input list. It is worth mentioning that the 3D-TCN-SE model is trained with the entire brain images for comparison reasons and not with just the ipsilateral hemisphere, as in the original paper [44].

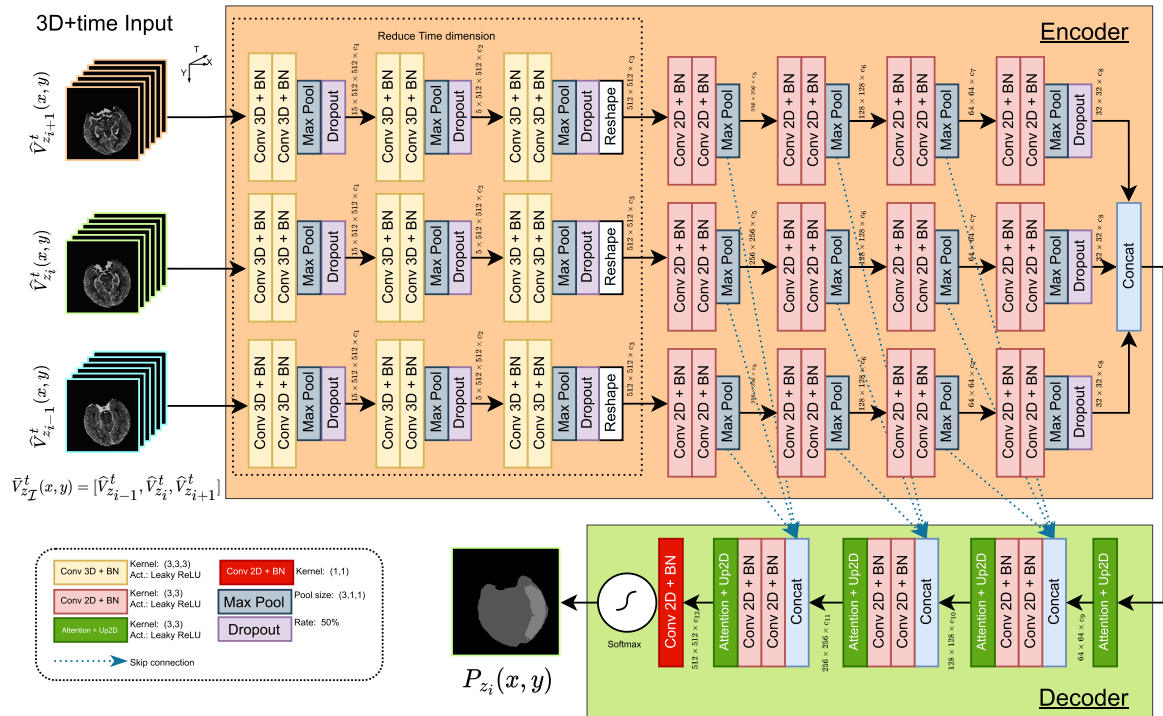


FIGURE 6. Illustration of the $3D+time$ mJ -Net model. The list of $2D+time$ input $\hat{V}_{z_{\mathcal{I}}}^t(x, y) = [\hat{V}_{z_{i-1}}^t, \hat{V}_{z_i}^t, \hat{V}_{z_{i+1}}^t]$ is trained in parallel, where $z_{\mathcal{I}} = \{z_{i-1}, z_i, z_{i+1}\}$. The output is a $2D$ image $P_{z_i}(x, y)$. The first max-pooling layer of each block in the convolution section has a pool size of $(2,1,1)$ to reduce the first dimension by a factor of 2. The second max-pooling layer uses a pool size of $(3,1,1)$, while the third has a pool size of $(5,1,1)$. The selection of these pool sizes is due to reducing the time dimension. The remaining max-pooling layers have a pool size of $(2,1,1)$. The Attention layers utilize a kernel of dimension 3 and a Leaky ReLU activation function. The $2D$ Upsampling layers have an upsampling factor of 2. The last convolution layer has a kernel of 1 and a Softmax activation function to produce a probability score for every class.

B. PROPOSED 4D METHODS

The proposed methods² adopt the entire 4D CTP scan as input to segment the ischemic regions (core and penumbra) in patients suspected of AIS. The main difference lies in how the 4D input is processed. The $3D$ -TCN is based on a $2D$ -TCN [43], modified to receive a list of 3D input volumes. The $3D+time$ mJ -Net inputs a list of $2D+time$ brain volumes from a CTP dataset, while the $4D$ mJ -Net uses the entire 4D structure of a CTP dataset as input. Fig. 5 compares these architectures with their respective inputs.

1) APPROACH 5: 3D TEMPORAL CONVOLUTIONAL NETWORK

We extend the architecture proposed by Amador et al. [43] for our application to exploit further the information in the depth dimension. In the remainder of the paper, we call our architecture $3D$ -TCN. The main differences between the proposed $3D$ -TCN and the $3D$ -TCN-SE (Sec. V-A4) rely on the usage of a 3D encoder for each input element, instead of a 3D single encoder, plus the possibility to segment both core and penumbra regions, in comparison with segmenting only the core areas.

²The code is publicly available at the following link <https://github.com/Biomedical-Data-Analysis-Laboratory/4D-mJ-Net>

The 4D CTP scans for the $3D$ -TCN are all re-sampled to 1 second per time point. As described in Sec. V-A3, the $3D$ -TCN architecture feeds each element of the input list $\hat{V}_{z_{\mathcal{I}}}^t$ at time point t_j to a specific 3D encoder $E_{3D-TCN}^{t_j}$ to extract low-level features. Each $E_{3D-TCN}^{t_j}$ encoder returns a $(4 \times 4 \times C)$ feature vector, where C corresponds to the number of channels. Each feature vector is merged to create a single input $ETOT_{3D-TCN} = [E_{3D-TCN}^{t_j} | \forall t_j \in \mathcal{I}]$. The $ETOT_{3D-TCN}$ is used in the TCN, which generates a one-dimensional vector O_{3D-TCN} of 64 elements. The TCN's output O_{3D-TCN} is then given in input to the decoder to create the final predicted 2D image $P_{z_i}(x, y)$ of a brain slice z_i at index i

2) APPROACH 6: 3D+TIME MJ-NET

We propose a model called $3D+time$ mJ -Net, an extension of the work of Tomasetti et al. [53]. The proposed model inputs a list of $2D+time$ matrices; thus, the dimension of this input can be defined as $3D+time$. The input and output are presented in Table 3, whereas a visual example of the input for the model is given in Fig. 3. Each element of the input list coincides with a possible input for the mJ -Net (details in Sec. V-A2). $\hat{V}_{z_{\mathcal{I}}}^t$ consists of a list of $2D+time$ volumes, where $z_{\mathcal{I}} = \{z_{i-1}, z_i, z_{i+1}\}$ is a set of brain slices containing the i th slice z_i analyzed and its neighboring slices z_{i-1} and z_{i+1} . In case the index i corresponds to the first (or last) brain slice, $\hat{V}_{z_{i-1}}^t$ (and

equivalently $\widehat{V}_{z_{i+1}}^t$ is set equal to $\widehat{V}_{z_i}^t$. Every 2D+time volume from the input list $\widehat{V}_{z_i}^t(x, y)$ is trained separately in the model through a series of encoders ($E_{z_{i-1}}, E_{z_i}, E_{z_{i+1}}$) composed of 3D-Conv and 2D-Conv layers. Each section is independent of the other: the convolution layers have no shared weights.

Fig. 6 illustrates the model's architecture. Attention layers [64] and 2D Upsampling layers were implemented in the decoder section. Attention layers benefit the architecture by focusing on target structures and help increase the segmentation performances. With the sole exception of the last convolution layer, each convolution layer uses a kernel of dimension 3 and a Leaky ReLU activation function [65] with $\alpha = 1/3$.

3) APPROACH 7: 4D MJ-NET

We propose another model called *4D mJ-Net*. We introduce this method to avoid the three paths to process the 4D data presented in the previous architecture (Sec. V-B2). We still use a sliding window technique over the depth dimension to limit the amount of input data fed to the model, using three consecutive brain slices simultaneously. Using a neighborhoods operation is a common image processing technique to include, in the computation, information from a additional dimension. Thus, adopting three consecutive brain slices enhances the model's ability to capture the information from the depth dimension effectively. Like the *3D+time mJ-Net* model, also this approach is an extension of the work of Tomasetti et al. [53]. Information on this approach is given in Table 3. The 4D input tensor $\widehat{V}_{z_{\mathcal{I}}}^t$ contains both the time dimension and the neighboring slices of the i th brain slice. The $\widehat{V}_{z_{\mathcal{I}}}^t$ is a concatenation of a 2D+time volume $\widehat{V}_{z_i}^t$ of a brain slice z_i at index i over all the time points t together with its neighbouring 2D+time volumes $\widehat{V}_{z_{i-1}}^t, \widehat{V}_{z_{i+1}}^t$. This model can be considered an early-fusion approach since the 4D input tensors $\widehat{V}_{z_{i-1}}^t, \widehat{V}_{z_i}^t, \widehat{V}_{z_{i+1}}^t$ are concatenated before being fed to the encoder's model.

The proposed *4D mJ-Net* model is a combination of both *3D+time mJ-Net* (Sec. V-B2) and *mJ-Net* (Sec. V-A2). The proposed approach uses the same input type that the *3D+time mJ-Net* exploits. However, rather than a list of 2D+time volumes, the model concatenates the input into a single 4D tensor of dimensions $(X \times Y \times Z \times T)$.

Unlike 1D, 2D, and 3D Convolution layers, 4D Convolution layers are not available in public DL frameworks (i.e., Keras³ or PyTorch⁴). Thus, for this model, we implemented a novel 4D-Conv layer (details in Sec. IV-C), which uses the convolutional layers defined in the public DL frameworks to replicate a 4D convolution operation.

The architecture of the *4D mJ-Net* is displayed in Fig. 7. No Attention layers [64] were included in these models due to a considerable performance decline. The output of the 4D-Conv layers is a tensor where the temporal dimension has been squeezed and reduced; information are extrapolated

from the temporal dimension. Thus, the output resulting from the 4D-Conv layers contains only three dimensions $(X \times Y \times Z)$ plus the channel dimension. 3D-Conv layers are implemented to reduce the depth dimension Z and produce a 2D vector $(X \times Y)$ plus the channel dimension.

A weighted categorical cross-entropy (WCC) loss [66] was the loss function implemented for this method. The loss can be written as:

$$\text{WCC}(x, y) = \sum_c \sum_i^{M \times N} (y_{i,c} \log x_{i,c}) \cdot (w_{i,c} y_{i,c}),$$

where $w_{i,c}$ corresponds to the weight of the i th pixel for a class $c \in \mathcal{C}$; $x_{i,c}$ is the i th predicted pixel, and $y_{i,c}$ is the corresponding ground truth pixel.

TABLE 3. Summary of the approaches. C-Time stands for computational time per brain slice in seconds. The C-Time values are expressed as mean \pm standard deviation.

Approach	Input	Output	Loss	C-Time (s)
Approach 1 [54]: Multi-input PMs	$\text{PMs}_{z_i} = [\text{CBF}_{z_i}, \text{CBV}_{z_i}, \text{TMAX}_{z_i}, \text{TTP}_{z_i}, \text{MIP}_{z_i}]$	$P_{z_i}(x, y)$	FTL [60]	0.07 ± 0.01
Approach 2 [53]: <i>mJ-Net</i>	$\widehat{V}_{z_i}^t(x, y) = \varphi(\widehat{I}_{z_i}^t(x, y) \forall t_j \in t)$	$P_{z_i}(x, y)$	SDCL [61]	0.19 ± 0.02
Approach 3 [43]: 2D-TCN	$\widehat{V}_{z_i}^t = [\widehat{I}_{z_i}^t \forall t_j \in t]$ of 2D images $\widehat{I}_{z_i}^t$	$P_{z_i}(x, y)$	DCL	0.40 ± 0.07
Approach 4 [44]: 3D-TCN-SE	$\widehat{V}_{z_{\mathcal{I}}}^t = [\widehat{V}_{z_i}^t \forall t_j \in t]$	$P_{z_i}(x, y)$	SDCL [61]	0.54 ± 0.06
Approach 5: 3D-TCN	$\widehat{V}_{z_{\mathcal{I}}}^t = [\widehat{V}_{z_i}^t \forall t_j \in t]$	$P_{z_i}(x, y)$	SDCL [61]	0.61 ± 0.02
Approach 6: 3D+time <i>mJ-Net</i>	$\widehat{V}_{z_{\mathcal{I}}}^t(x, y) = [\widehat{V}_{z_{i-1}}^t, \widehat{V}_{z_i}^t, \widehat{V}_{z_{i+1}}^t]$	$P_{z_i}(x, y)$	SDCL [61]	0.65 ± 0.02
Approach 7: <i>4D mJ-Net</i>	$\widehat{V}_{z_{\mathcal{I}}}^t = \varphi(\widehat{V}_{z_{i-1}}^t, \widehat{V}_{z_i}^t, \widehat{V}_{z_{i+1}}^t)$	$P_{z_i}(x, y)$	WCC [66]	0.38 ± 0.03

C. IMPLEMENTATION DETAILS

Table 3 provides information about all the methods. All the methods mentioned in Sec. V-A and Sec. V-B utilize Adam as the optimizer [67] with a learning rate of 0.0003 and a step-based decay rate of 0.95 every ten epochs. The batch size is set to 2. An early stopping function is called if there is no decrement in the validation loss after 25 epochs. During training, L1 and L2 regularizations are applied in the kernels, plus a max norm constraint is also used in the kernel and bias weights. The average computational time for all the brain slices during inference phase is also highlighted in Table 3. All experiments were implemented in Python using Keras (2.3.1) with Tensorflow as the backend and trained using an NVIDIA Tesla V100 GPU (32 GB memory).

VI. EXPERIMENTS AND RESULTS

We assess the proposed methods⁵ on a local dataset of CTP scans from 152 patients (Sec. III). All experiments are performed with the same training set and evaluated over the validation set (details in Table 1). The test set is used only to make predictions with the best models with CTP scans that the methods have not seen before. Since the Non-LVO group has smaller ischemic areas than the LVO patients, we set a

³<https://keras.io/>

⁴<https://pytorch.org/>

⁵The code is publicly available at the following link: <https://github.com/Biomedical-Data-Analysis-Laboratory/4D-mJ-Net>

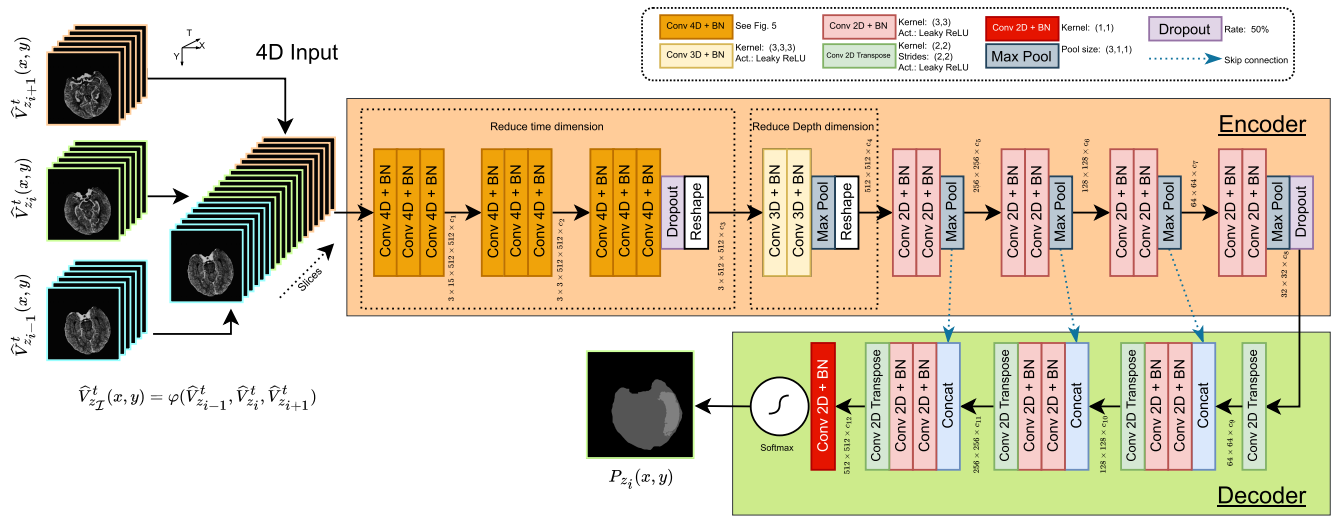


FIGURE 7. Illustration of the 4D *mJ-Net* architecture. The 4D input $\hat{V}_{z_I}^t = \varphi(\hat{V}_{z_I-1}^t, \hat{V}_{z_I}^t, \hat{V}_{z_I+1}^t)$ is the concatenation of a 2D+time volume $\hat{V}_{z_I}^t$ of a brain slice z_I at index i over all the time points t plus its neighbouring brain slice volumes ($\hat{V}_{z_I-1}^t, \hat{V}_{z_I+1}^t$). Two MonteCarlo dropout layers [63] are added at the end of the 4D and 2D Convolution blocks. The rate was set to 50%. These layers were added to reduce uncertainties in the final predictions. Additional details are given in the supplemental material. The last convolution layer has a kernel of (1×1) and a Softmax activation function to produce a probability score for every class.

TABLE 4. Experiment results for the validation set. Values in bold exhibit the best results for each column and each class. Mean results plus standard deviation for Dice Coefficient (DC), Hausdorff Distance (HD), and ΔV are presented. Results are for the penumbra and core areas divided by the distinct patient groups (LVO, Non-LVO, WIS, and All). Note that for the DC, higher values are better (\uparrow), while for HD and ΔV , lower values are preferable (\downarrow).

Method	DC \uparrow			HD (mm) \downarrow			ΔV (ml) \downarrow			
	LVO	Non-LVO	All	LVO	Non-LVO	All	LVO	Non-LVO	WIS	All
Penumbra										
<i>Multi-input PMs</i> [54]	0.70 \pm 0.1	0.27 \pm 0.3	0.47 \pm 0.3	2.9 \pm 0.4	1.4 \pm 0.7	2.0 \pm 0.8	27.0 \pm 28.6	10.0 \pm 15.5	9.8 \pm 8.1	19.0 \pm 24.2
<i>mJ-Net</i> [53]	0.66 \pm 0.2	0.39 \pm 0.3	0.50 \pm 0.3	2.9 \pm 0.5	2.6 \pm 0.6	2.7 \pm 0.7	25.5 \pm 20.0	24.7 \pm 29.2	45.5 \pm 39.1	27.2 \pm 26.0
<i>2D-TCN</i> [43]	0.12 \pm 0.1	0.02 \pm 0.0	0.07 \pm 0.1	4.1 \pm 0.5	3.8 \pm 0.6	4.0 \pm 0.6	81.3 \pm 65.6	80.6 \pm 57.8	131.6 \pm 93.1	86.0 \pm 66.5
<i>3D-TCN-SE</i> [44]	0.25 \pm 0.1	0.05 \pm 0.1	0.15 \pm 0.1	6.2 \pm 0.5	6.7 \pm 0.4	6.4 \pm 0.5	497.9 \pm 157.1	559.3 \pm 90.4	624.6 \pm 118.7	533.1 \pm 137.3
<i>3D-TCN</i>	0.23 \pm 0.1	0.04 \pm 0.1	0.14 \pm 0.1	4.3 \pm 0.4	4.4 \pm 0.5	4.3 \pm 0.5	85.3 \pm 64.0	142.7 \pm 51.4	164.2 \pm 43.8	114.2 \pm 65.3
<i>3D+time mJ-Net</i>	0.70 \pm 0.1	0.42 \pm 0.3	0.53 \pm 0.3	2.6 \pm 0.6	1.9 \pm 0.8	2.2 \pm 0.9	35.1 \pm 36.1	18.3 \pm 26.3	2.7 \pm 2.6	25.7 \pm 32.4
<i>4D mJ-Net</i>	0.66 \pm 0.1	0.44 \pm 0.3	0.51 \pm 0.3	2.3 \pm 0.6	1.3 \pm 0.7	1.7 \pm 1.0	41.4 \pm 37.2	6.1 \pm 6.3	0.0 \pm 0.0	24.3 \pm 32.9
Core										
<i>Multi-input PMs</i> [54]	0.37 \pm 0.3	0.21 \pm 0.3	0.28 \pm 0.3	1.2 \pm 0.8	0.4 \pm 0.4	0.8 \pm 0.8	9.4 \pm 20.3	0.8 \pm 1.3	0.5 \pm 0.5	5.3 \pm 15.3
<i>mJ-Net</i> [53]	0.27 \pm 0.2	0.21 \pm 0.2	0.22 \pm 0.2	1.5 \pm 0.7	0.8 \pm 0.6	1.2 \pm 0.6	5.5 \pm 4.9	1.0 \pm 1.2	1.0 \pm 1.1	3.4 \pm 4.3
<i>2D-TCN</i> [43]	0.02 \pm 0.0	0.01 \pm 0.0	0.01 \pm 0.0	1.9 \pm 0.7	1.5 \pm 0.6	1.7 \pm 0.7	11.8 \pm 13.3	8.1 \pm 8.2	11.0 \pm 11.2	10.3 \pm 11.4
<i>3D-TCN-SE</i> [44]	0.00 \pm 0.0	0.00 \pm 0.0	0.00 \pm 0.0	1.2 \pm 0.9	0.4 \pm 0.4	0.8 \pm 0.8	12.7 \pm 15.6	1.9 \pm 2.8	0.0 \pm 0.0	7.5 \pm 12.8
<i>3D-TCN</i>	0.02 \pm 0.0	0.01 \pm 0.0	0.01 \pm 0.0	1.4 \pm 0.8	0.8 \pm 0.4	1.1 \pm 0.7	12.0 \pm 14.3	1.9 \pm 2.1	2.4 \pm 1.9	7.3 \pm 11.6
<i>3D+time mJ-Net</i>	0.21 \pm 0.2	0.12 \pm 0.2	0.16 \pm 0.4	1.1 \pm 0.7	0.4 \pm 0.4	0.7 \pm 0.7	8.1 \pm 10.6	1.3 \pm 1.6	0.0 \pm 0.0	4.8 \pm 8.5
<i>4D mJ-Net</i>	0.29 \pm 0.2	0.21 \pm 0.2	0.23 \pm 0.2	1.6 \pm 0.9	0.5 \pm 0.4	1.0 \pm 0.9	25.9 \pm 37.0	1.4 \pm 2.2	0.0 \pm 0.0	14.3 \pm 29.6

higher penalty for every misclassification of penumbra and core classes for this sub-group during training.

A. EVALUATION METRICS

Three evaluation metrics are used to assess the various experiments' models. The Dice Coefficient (DC), the Hausdorff Distance (HD) [68], and the absolute difference in the volumes (ΔV). We employ the DC to compare the model predictions with the ground truth segmentations. The DC between two segmentations x and y is given by the following equation:

$$DC(x, y) = 2 \frac{|x \cap y|}{|x| + |y|}$$

where the range for the DC is $[0, 1]$; thus a $DC(x, y) = 1$ corresponds to a perfect match between the prediction x and ground truth y segmentations.

The HD measures how two subsets (\mathcal{A}, \mathcal{B}) are distant from each other, and it is formulated as follows:

$$HD(\mathcal{A}, \mathcal{B}) = \max \{h(\mathcal{A}, \mathcal{B}), h(\mathcal{B}, \mathcal{A})\},$$

where $h(\mathcal{A}, \mathcal{B}) = \max_{a \in \mathcal{A}} \min_{b \in \mathcal{B}} \|a - b\|$. The range value for the HD is $[0, \infty]$.

The absolute difference in the volumes ΔV between the prediction volume V_x and the ground truth volume V_y can be expressed as:

$$\Delta V(V_y, V_x) = |V_y - V_x|$$

The range for ΔV is $[0, \infty]$, and $\Delta V(V_y, V_x) = 0$ represents a perfect match between the two volumes. The $\Delta V(V_y, V_x)$ is an essential evaluation metric for the WIS group due to the lack of ground truth segmentations in this group. The other metrics are not suitable for understanding how the predictions will be since the ground truth will always be empty.

The best scenario for a model is to produce high DC with low HD and ΔV : this implies a strong correlation between the predicted areas and the ground truth regions. If the results show high ΔV (or HD) with low DC, an over-segmentation of the ischemic areas is perceived. On the other hand, promising outcomes of ΔV (or HD) with mediocre DC results imply an under-segmentation of the predicted regions.

TABLE 5. Ablation study for the 4D *mJ-Net* model showing how various pre-processing steps (HE, γ , z) and re-sampling (ψ) affect the Dice Coefficient (DC) for the validation results. Penumbra and core DC scores are shown for all the classes together. Note that for the DC, higher values are better (\uparrow).

Ablation Setting					DC \uparrow			
					LVO		Non-LVO	
HE	γ	z	ψ	Penumbra	Core	Penumbra	Core	
-	-	-	-	0.42 \pm 0.2	0.25 \pm 0.2	0.20 \pm 0.2	0.16 \pm 0.2	
✓	-	-	-	0.32 \pm 0.3	0.24 \pm 0.2	0.13 \pm 0.2	0.16 \pm 0.3	
-	✓	-	-	0.00 \pm 0.0	0.07 \pm 0.1	0.00 \pm 0.0	0.06 \pm 0.1	
✓	✓	-	-	0.48 \pm 0.2	0.28 \pm 0.2	0.24 \pm 0.2	0.20 \pm 0.3	
-	-	✓	-	0.01 \pm 0.0	0.14 \pm 0.2	0.01 \pm 0.0	0.08 \pm 0.1	
✓	-	✓	-	0.53 \pm 0.2	0.22 \pm 0.2	0.35 \pm 0.3	0.12 \pm 0.2	
-	✓	✓	-	0.28 \pm 0.2	0.17 \pm 0.2	0.08 \pm 0.1	0.05 \pm 0.1	
✓	✓	✓	-	0.66\pm0.1	0.29\pm0.2	0.44\pm0.3	0.21\pm0.2	
-	-	-	✓	0.00 \pm 0.0	0.00 \pm 0.0	0.00 \pm 0.0	0.00 \pm 0.0	
✓	-	-	✓	0.07 \pm 0.1	0.26 \pm 0.2	0.01 \pm 0.0	0.05 \pm 0.1	
-	✓	-	✓	0.00 \pm 0.0	0.00 \pm 0.0	0.00 \pm 0.0	0.00 \pm 0.0	
✓	✓	-	✓	0.41 \pm 0.3	0.29 \pm 0.2	0.11 \pm 0.2	0.12 \pm 0.2	
-	-	✓	✓	0.00 \pm 0.0	0.00 \pm 0.0	0.00 \pm 0.0	0.00 \pm 0.0	
✓	-	✓	✓	0.56 \pm 0.2	0.24 \pm 0.2	0.37 \pm 0.3	0.18 \pm 0.2	
-	✓	✓	✓	0.00 \pm 0.0	0.00 \pm 0.0	0.00 \pm 0.0	0.00 \pm 0.0	
✓	✓	✓	✓	0.59 \pm 0.2	0.29\pm0.2	0.40 \pm 0.3	0.20 \pm 0.2	

B. COMPARISON WITH OTHER METHODS

The proposed 3D+time *mJ-Net*, 4D *mJ-Net*, and 3D-TCN methods are compared with alternative models: the 2D-TCN [43], the *mJ-Net* [53], the 3D-TCN-SE [44], and the Multi-input PMs [54].

Table 4 presents the evaluation metrics' results over the validation set. Results are presented for each group distinctly (LVO, Non-LVO, WIS, and all) to highlight the strengths and weaknesses of each model over the various groups composing the dataset. An extensive number of experiments are performed for all the analyzed models. However, to present a fair comparison among the various models, we only introduce the methods with a combination of parameters that yield the best results, omitting the other combinations tested during experiments. Qualitative comparison results of random brain slices extracted from the validation set are provided in Fig. 8.

C. ABLATION STUDY

To demonstrate the effects of the pre-processing steps (Sec. IV-B), we conduct an ablation study on the 4D *mJ-Net* architecture. Moreover, we re-sampled the CTP scans

to handle the irregular temporal dimension and studied the effect of using re-sampled scans during the model's training. Different CT scan vendors have different imaging acquisition protocols; thus, re-sampling the scans to a fixed time-sampling rate is a reasonable step to increase the versatility and usability across hospitals. DC is illustrated in Table 5, showing performances of the network for all the groups trained with the datasets using different types of pre-processing steps and re-sampled scans. The study aims to systematically analyze the contribution of each pre-processing step toward improving the overall results. We begin by defining a baseline configuration consisting of the raw input images without pre-processing (first row in Table 5). Subsequently, we incrementally introduce and evaluate individual pre-processing steps, such as histogram equalization (HE), gamma correction (γ), and z-score (z).

D. INTER-OBSERVER VARIABILITY

Two expert neuroradiologists (NR₁, NR₂) manually annotated the scans of 33 randomly selected patients: 19 from the LVO group, 11 from the Non-LVO, and 3 from the WIS subset. The manual annotation images were generated using the same criteria endorsed for creating the ground truth images, as explained in Sec. III-A. An investigation of the inter-observer variability between NR₁, NR₂, and the two best-proposed models is presented in Table 6.

VII. DISCUSSION

Early detection and intervention in AIS patients are of vital importance [69], [70], [71]. In this study, we have proposed different architectures to utilize the 4D CTP input to use the spatio-temporal information better than in existing approaches. We suggest expanding the *mJ-Net* and showing two ways of segmenting ischemic areas in patients suspected of AIS. In addition, we expand another method (3D-TCN) for comparison reasons. We use the entire raw 4D CTP data and feed different combinations as input to our proposed approaches to prevent possible loss of spatio-temporal information. Studying the data as an independent volume and neglecting its spatio-temporal nature can lead to the loss of relevant information. All proposed approaches return a series of 2D segmented images as output, later stacked together to produce a 3D volume. Returning a list of 2D images as output is less computationally expensive and less memory intensive than directly returning 3D volumetric data as output [72].

Few studies have adopted 4D datasets in DNN models to detect ischemic lesions in patients affected by a stroke [9], [43], [44], [45]. This is rooted in the high computational complexity of 4D data and the lack of ground truth for the whole set. The limitations that these approaches encounter are as follows: 1) datasets used for the training and evaluation take into account only a subset of the entire population; 2) segmentations are only performed on the core areas, excluding penumbra regions; 3) ground truth images derived

TABLE 6. Inter-observer variability results for test set. Values are presented for the two best-proposed architectures (*3D+time mJ-Net*, *4D mJ-Net*) in relation to manual annotations generated separately by two expert neuroradiologists (NR_1 , NR_2) over the test set. An investigation of the inter-observer variability between NR_1 and NR_2 is performed (last row of each class). Note that for the DC, higher values are better (\uparrow), while for HD and ΔV , lower values are preferable (\downarrow).

Method	DC \uparrow			HD (mm) \downarrow			ΔV (ml) \downarrow			
	LVO	Non-LVO	All	LVO	Non-LVO	All	LVO	Non-LVO	WIS	All
Penumbra										
<i>3D+time mJ-Net</i> vs (NR_1 vs NR_2)	0.70 \pm 0.1	0.32 \pm 0.3	0.51 \pm 0.3	2.6 \pm 0.5	1.5 \pm 0.6	2.1 \pm 0.9	36.7 \pm 36.4	6.5 \pm 4.9	10.1 \pm 3.1	24.2 \pm 31.1
<i>4D mJ-Net</i> vs (NR_1 vs NR_2)	0.67 \pm 0.1	0.25 \pm 0.3	0.47 \pm 0.3	2.4 \pm 0.5	0.9 \pm 0.5	1.7 \pm 1.0	34.1 \pm 30.6	5.3\pm6.6	0.0\pm0.0	21.4 \pm 27.7
NR_1 vs NR_2	0.78\pm0.1	0.65\pm0.2	0.67\pm0.2	2.2\pm0.4	0.8\pm0.6	1.6\pm1.0	33.3\pm27.7	5.5 \pm 9.2	0.0\pm0.0	21.0\pm25.9
Core										
<i>3D+time mJ-Net</i> vs (NR_1 vs NR_2)	0.19 \pm 0.2	0.01 \pm 0.0	0.12 \pm 0.2	1.6 \pm 0.8	0.3 \pm 0.4	1.0 \pm 1.0	14.6 \pm 18.2	0.9 \pm 2.3	0.0\pm0.0	8.7 \pm 15.4
<i>4D mJ-Net</i> vs (NR_1 vs NR_2)	0.28 \pm 0.2	0.03 \pm 0.1	0.18 \pm 0.2	1.8 \pm 0.8	0.3 \pm 0.5	1.1 \pm 1.0	21.2 \pm 31.7	1.8 \pm 4.5	0.0\pm0.0	12.8 \pm 25.9
NR_1 vs NR_2	0.44\pm0.2	0.15\pm0.2	0.30\pm0.3	1.4\pm0.6	0.2\pm0.4	0.9\pm0.8	5.6\pm4.3	0.7\pm1.9	0.0\pm0.0	3.5\pm4.2

from follow-up DWI or NCCT present some limitations [47], [48].

To our knowledge, this is the first study using 4D CTP data to segment both the ischemic regions, penumbra and core. Additionally, we include data from all patients, regardless of stroke severity, to train our models. Rather than entrusting ground truth images from follow-up DWI or NCCT studies that are usually taken 24 hours or some days after stroke onset, our proposed methods were trained with ground truth images obtained from the CTP captured at admission, including PMs, and follow-up scans (Sec. III-A).

We use three evaluation metrics to assess the models' performances: DC, HD, and ΔV compared with our previously developed algorithms and other state-of-the-art algorithms. Results in Table 4 demonstrate that increasing the input dimension benefits achieving more precise segmentation, especially for the Non-LVO and WIS groups, regardless of the class. Thus, when a smaller portion of the brain is affected, the whole dataset's usage helps achieve better segmentation results. The ablation study (Table 5) shows how including the pre-processing steps and not re-sampling the CTP scans helped improve the overall segmentation performances. It is possible to evince that by combining multiple pre-processing steps, the overall segmentation performances increased, regardless of whether using or not the re-sampling technique. Adopting the HE step with other pre-processing steps is beneficial for the AIS segmentations; the absence of the HE step decreases the overall performance. It is worth mentioning that using the pre-processing steps and re-sampled CTP scans yields the second-best overall results (last row in Table 5), establishing the validity of the pre-processing sequence.

Additionally, we evaluated the computational time for predicting each brain slice (Table 3), highlighting how fast each method performs. From the results in Table 3, it is possible to notice that the *Multi-input PMs* architecture has the lowest computational cost due to the usage of PMs instead of CTP scans as input. Among the methods utilizing CTP scans as input, the *mJ-Net* has the best computational time

since it consists of a more straightforward structure compared to the other models.

Visual results of random validation brain slices are shown in Fig. 8, where we can see that our proposed approaches (*3D+time mJ-Net*, *4D mJ-Net*) are less prone to over-segment, especially in the Non-LVO and WIS groups. It is reported that LVO cases are less common compared to Non-LVO. On average, LVOs are estimated to represent around 30% of all AIS cases [73]. Thus, a neural network that can accurately segment patients in the Non-LVO group can be valuable in a real-life scenario. Nonetheless, patients with LVO represent a clinically significant proportion of patients presenting with AIS, especially considering the grim natural course of the disease.

The results presented in Table 4 indicate that all *mJ-Net* models have improved where the input data dimension has increased, regardless of the patients' group. Fig. 8 shows that using 2D+time input for the *mJ-Net* [53] led to over-segmentation of penumbra class in separate brain tissue sections, brain slice 4-6. The visual results for the Non-LVO and WIS groups highlight the limitations of this model: the over-segmentation of the penumbra regions might affect the usage in a real-life scenario, and an overestimation of the penumbra area can generate uncertainties for treatment decisions.

Adding depth as an extra dimension to the input of models (*3D+time mJ-Net* and *4D mJ-Net*) determines an increment in the performances for both classes in the three patient groups. A significant increase is noticeable for the DC metric in the Non-LVO group, regardless of the class. An essential difference between these two architectures is how they exploit their structures' input. The *3D+time mJ-Net* is considered a late-fusion approach as the data sources are used independently and fused close to decision-making. Statistical results presented for the *3D+time mJ-Net* show promising general performances for the LVO group. However, an underestimation of the core class, regardless of the patient group, can be noticed from the visual results in




FIGURE 8. Qualitative comparisons for the tested models. The brain slices are taken randomly from distinct patients from the validation set, divided by group (left to right). Each row represents the results for each model involved in the study. The numbers on the bottom sides of every image represent the DC for penumbra (in light blue) and core (in green).

Fig. 8 and the low HD metric. Nevertheless, *3D+time mJ-Net* achieved the best HD for the core class in all the groups. The *3D+time mJ-Net* can precisely segment ischemic regions with large areas, as shown by the first three slices in Fig. 8. This can also be manifested in the high DC score achieved for the LVO group for both classes (Table 4).

The *4D mJ-Net* network has learned to precisely segment the ischemic regions even without re-sampling the CTP scans, as shown in Table 5. The model fuses the data before they are fed to the network. Visual results in Fig. 8 and values in Table 4 indicate that the *4D mJ-Net* model segments the penumbra class more precisely compared to the other approaches that use raw CTP as their input. This promising performance follows in all patient groups. The *4D mJ-Net* achieved the highest DC metric for core and penumbra regions in patients with Non-LVO. This approach gives the best HD for penumbra in all the groups. The *4D mJ-Net* showed high precision in detecting small ischemic areas, as shown in sample brain slices 3 to 5 in Fig. 8. Furthermore, the *4D mJ-Net* model can correctly predict no ischemic regions in WIS patients, as demonstrated by the results for the ΔV in the WIS group. However, it over-segments the core class in patients with LVO. This means that including the complete spatio-temporal information of the data and following an early fusion approach leads to better prediction in Non-LVO and WIS groups, where small areas are of interest.

Models based on TCN generally showed poor results statistically in Table 4 and visually in Fig. 8. They extremely over-segment the penumbra class and poorly segment the core class. The original *2D-TCN* and *3D-TCN-SE* were designed to segment only one class, the ischemic core. This can explain the poor performance of segmenting the two classes. Besides, in Amador et al. [44] (*3D-TCN-SE*), the model was trained to use only the ipsilateral hemisphere. For a fair comparison, the model's training was done over both hemispheres, which can cause over-segmentation in penumbra regions.

As the name indicates, the *Multi-input PMs* model [54] takes parametric maps and pre-processed data obtained from CTP scans. The experiment results of this model show a high DC value for the penumbra class in the LVO group, as also seen in the first three brain slices of Fig. 8. This highlights that this method presents satisfactory results for large ischemic areas. However, when the region's volume is small or vacant, the predictions are not optimal: see brain slices 4 and 6 in Fig. 8. Although HD and ΔV are encouraging, DC values show under-segmentation in the core and penumbra classes for the Non-LVO set. Using PMs derived from CTP scans limits the machine to only learn from specific pre-processed information.

The inter-observer variability results, highlighted in Table 6, show promising outcomes for the proposed methods with the results achieved by the two expert neuroradiologists (NR_1 , NR_2). Similar statistic values can be observed between NR_1 vs. NR_2 and the *4D mJ-Net* for the penumbra class

in connection with the LVO group. The same results as the neuroradiologists were achieved by the *4D mJ-Net* for the ΔV in the WIS group. The proposed *3D+time mJ-Net* model produces higher results for the DC compared to the *4D mJ-Net* in association with the penumbra class. However, the results for the core regions could be more satisfactory. The inter-observer variability outcomes for the HD and ΔV highlight substantial similarity among the proposed approaches and the neuroradiologists, except for the core class connected with the LVO group. The difference can be due to an over-segmentation of this particular class, which can be highly complex to detect for the models, considering its small size.

To summarize the discussion, the proposed *3D+time mJ-Net* model can precisely segment large ischemic penumbra regions while presenting an underestimation of the core class. The proposed *4D mJ-Net* network can segment small ischemic areas with high precision and predict the absence of ischemic regions; nevertheless, it presents an overestimation of the core class in LVO patients.

A. COMMON LIMITATIONS

All the assessed approaches have faced general limitations. The images used during the training of each model are from CT scanners of the same vendor. This causes a lack of diversity in the data. The annotations used as ground truth surround the essential ischemic regions (penumbra and core) but do not represent the areas perfectly. They might leave out small parts of the core spread into the penumbra tissue and details of the penumbra misclassified as healthy brain tissue [53]. Additional qualitative results for all the methods, showing their worst segmentation results, are presented in the supplemental material.

VIII. CONCLUSION

Fast and precise diagnosis and treatment are of vital importance in AIS patients. In this paper, we proposed to use 4D CTP as input to extract spatio-temporal information for segmenting core and penumbra areas in patients with AIS. This is presented primarily by expanding the *mJ-Net* in two ways (*3D+time mJ-Net* and *4D mJ-Net*). Furthermore, we introduced a novel 4D-Conv layer to exploit spatio-temporal information. Two of our approaches (*3D+time mJ-Net* and *4D mJ-Net*) achieved promising results for all the classes involved. The *3D+time mJ-Net* can precisely delineate large ischemic penumbra areas while underestimating the core class. Our best network (*4D mJ-Net*) can correctly segment penumbra regions, regardless of patient groups, with a 0.53 DC score on average. However, with an average of 0.23 DC score, it overestimates the core class for the LVO group.

We used the entire 4D CTP dataset of all patients and compared models using different input types. We demonstrated that relying only on images derived from the CTP scans (i.e., PMs) or on a restricted number of dimensions (i.e., 2D, 2D+time, 3D) limits the prediction accuracy in DNN-based

approaches. Moreover, we segmented both penumbra and core regions in ischemic brain tissue since an accurate and fast understanding of both is essential for quick treatment decisions in AIS.

Further studies with larger datasets, including images from different vendors and various acquisition parameters, are still needed to validate our methods. The ISLES18 dataset [74] can be used in future work; the dataset uses FIAs as ground truth labels, which is not in the scope of the architectures mentioned above; thus, some changes must be implemented for validating the methods. Due to complex and time-consuming work for manual annotations, further work on optimizing the segmentation using unsupervised neural networks is encouraged.

ACKNOWLEDGMENT

The authors are with the BioMedical Data Analysis Group (<https://www.uis.no/en/bmdlab>). They declare that they do not have any commercial or associative interest that represents a conflict of interest in connection with the work submitted. strings

REFERENCES

- [1] *Optimizing Brain Health Across the Life Course: Who Position Paper*, World Health Organization, Geneva, Switzerland, 2022.
- [2] S. Ojaghhiaghghi, S. S. Vahdati, A. Mikaeilpour, and A. Ramouz, "Comparison of neurological clinical manifestation in patients with hemorrhagic and ischemic stroke," *World J. Emergency Med.*, vol. 8, no. 1, p. 34, 2017.
- [3] J. Astrup, B. K. Siesjö, and L. Symon, "Thresholds in cerebral ischemia—the ischemic penumbra," *Stroke*, vol. 12, no. 6, pp. 723–725, 1981.
- [4] European Stroke Organisation (ESO) Executive Committee, & ESO Writing Committee, "Guidelines for management of ischaemic stroke and transient ischaemic attack 2008," *Cerebrovascular Diseases*, vol. 25, no. 5, pp. 457–507, 2008.
- [5] B. C. V. Campbell, L. Weir, P. M. Desmond, H. T. H. Tu, P. J. Hand, B. Yan, G. A. Donnan, M. W. Parsons, and S. M. Davis, "CT perfusion improves diagnostic accuracy and confidence in acute ischaemic stroke," *J. Neurol., Neurosurg. Psychiatry*, vol. 84, no. 6, pp. 613–618, Jun. 2013.
- [6] A. Bivard, C. Levi, V. Krishnamurthy, J. Hislop-Jambrich, P. Salazar, B. Jackson, S. Davis, and M. Parsons, "Defining acute ischemic stroke tissue pathophysiology with whole brain CT perfusion," *J. Neuroradiol.*, vol. 41, no. 5, pp. 307–315, Dec. 2014.
- [7] K. D. Kurz, G. Ringstad, A. Odland, R. Advani, E. Farbu, and M. W. Kurz, "Radiological imaging in acute ischaemic stroke," *Eur. J. Neurol.*, vol. 23, no. S1, pp. 8–17, Jan. 2016.
- [8] A. M. McKinney, "Differences in CT perfusion maps generated by different commercial software: Quantitative analysis by using identical source data of acute stroke patients," *Yearbook Neurol. Neurosurg.*, vol. 2010, pp. 157–159, Jan. 2010.
- [9] M. Soltanpour, M. Yousefnezhad, R. Greiner, P. Boulanger, and B. Buck, "Using end-to-end imbalanced temporal deep Gaussian process (ITDGP) to predict acute ischemic stroke lesions," 2022, *arXiv:2211.09068*.
- [10] C. D. d'Esteire, M. E. Boesen, S. H. Ahn, P. Pordeli, M. Najm, P. Minhas, P. Davari, E. Fainardi, M. Rubiera, A. V. Khaw, A. Zini, R. Frayne, M. D. Hill, A. M. Demchuk, T. T. Sajobi, N. D. Forkert, M. Goyal, T. Y. Lee, and B. K. Menon, "Time-dependent computed tomographic perfusion thresholds for patients with acute ischemic stroke," *Stroke*, vol. 46, no. 12, pp. 3390–3397, Dec. 2015.
- [11] L. Lin, A. Bivard, V. Krishnamurthy, C. R. Levi, and M. W. Parsons, "Whole-brain CT perfusion to quantify acute ischemic penumbra and core," *Radiology*, vol. 279, no. 3, pp. 876–887, Jun. 2016.
- [12] F. McVerry, K. A. Dani, N. J. J. MacDougall, M. J. MacLeod, J. Wardlaw, and K. W. Muir, "Derivation and evaluation of thresholds for core and tissue at risk of infarction using CT perfusion," *J. Neuroimaging*, vol. 24, no. 6, pp. 562–568, Nov. 2014.
- [13] C. W. Cereda, S. Christensen, B. C. Campbell, N. K. Mishra, M. Mlynash, C. Levi, M. Straka, M. Wintermark, R. Bammer, G. W. Albers, M. W. Parsons, and M. G. Lansberg, "A benchmarking tool to evaluate computer tomography perfusion infarct core predictions against a DWI standard," *J. Cerebral Blood Flow Metabolism*, vol. 36, no. 10, pp. 1780–1789, Oct. 2016.
- [14] A. Nielsen, M. B. Hansen, A. Tietze, and K. Mouridsen, "Prediction of tissue outcome and assessment of treatment effect in acute ischemic stroke using deep learning," *Stroke*, vol. 49, no. 6, pp. 1394–1401, Jun. 2018.
- [15] A. Bivard, L. Churilov, H. Ma, C. Levi, B. Campbell, N. Yassi, A. Meretoja, H. Zhao, G. Sharma, C. Chen, S. Davis, G. Donnan, B. Yan, and M. Parsons, "Does variability in automated perfusion software outputs for acute ischemic stroke matter? Reanalysis of EXTEND perfusion imaging," *CNS Neurosci. Therapeutics*, vol. 28, no. 1, pp. 139–144, Jan. 2022.
- [16] C.-L. Chen, C.-C. Chen, W.-H. Yu, S.-H. Chen, Y.-C. Chang, T.-I. Hsu, M. Hsiao, C.-Y. Yeh, and C.-Y. Chen, "An annotation-free whole-slide training approach to pathological classification of lung cancer types using deep learning," *Nature Commun.*, vol. 12, no. 1, pp. 1–12, Feb. 2021.
- [17] N. Kanwal, S. Fuster, F. Khoraminia, T. C. M. Zuiverloon, C. Rong, and K. Engan, "Quantifying the effect of color processing on blood and damaged tissue detection in whole slide images," in *Proc. IEEE 14th Image, Video, Multidimensional Signal Process. Workshop (IVMSP)*, Jun. 2022, pp. 1–5.
- [18] S. Fuster, F. Khoraminia, U. Kiraz, N. Kanwal, V. Kvikstad, T. Eftestøl, T. C. M. Zuiverloon, E. A. M. Janssen, and K. Engan, "Invasive cancerous area detection in non-muscle invasive bladder cancer whole slide images," in *Proc. IEEE 14th Image, Video, Multidimensional Signal Process. Workshop (IVMSP)*, Jun. 2022, pp. 1–5.
- [19] N. Kanwal, F. Pérez-Bueno, A. Schmidt, K. Engan, and R. Molina, "The devil is in the details: Whole slide image acquisition and processing for artifacts detection, color variation, and data augmentation: A review," *IEEE Access*, vol. 10, pp. 58821–58844, 2022.
- [20] N. Kanwal, T. Eftestøl, F. Khoraminia, T. C. M. Zuiverloon, and K. Engan, "Vision transformers for small histological datasets learned through knowledge distillation," in *Advances in Knowledge Discovery and Data Mining*. Cham, Switzerland: Springer, 2023.
- [21] I. Chakraborty, A. Elgammal, and R. S. Burd, "Video based activity recognition in trauma resuscitation," in *Proc. 10th IEEE Int. Conf. Workshops Autom. Face Gesture Recognit. (FG)*, Apr. 2013, pp. 1–8.
- [22] Ø. Meinich-Bache, S. L. Austnes, K. Engan, I. Austvoll, T. Eftestøl, H. Myklebust, S. Kusulla, H. Kidanto, and H. Ersdal, "Activity recognition from newborn resuscitation videos," *IEEE J. Biomed. Health Informat.*, vol. 24, no. 11, pp. 3258–3267, Nov. 2020.
- [23] F. Patlar Akbulut, B. Ikitimur, and A. Akan, "Wearable sensor-based evaluation of psychosocial stress in patients with metabolic syndrome," *Artif. Intell. Med.*, vol. 104, Apr. 2020, Art. no. 101824.
- [24] T. Ahmad, S. T. H. Rizvi, and N. Kanwal, "Transforming spatio-temporal self-attention using action embedding for skeleton-based action recognition," *J. Vis. Commun. Image Represent.*, vol. 95, Sep. 2023, Art. no. 103892.
- [25] O. Amft and G. Tröster, "Recognition of dietary activity events using on-body sensors," *Artif. Intell. Med.*, vol. 42, no. 2, pp. 121–136, Feb. 2008.
- [26] O. Ronneberger, P. Fischer, and T. Brox, "U-Net: Convolutional networks for biomedical image segmentation," in *Proc. Int. Conf. Med. Image Comput. Comput.-Assist. Intervent.* Cham, Switzerland: Springer, 2015, pp. 234–241.
- [27] F. Isensee, P. F. Jaeger, S. A. A. Kohl, J. Petersen, and K. H. Maier-Hein, "nnU-Net: A self-configuring method for deep learning-based biomedical image segmentation," *Nature Methods*, vol. 18, no. 2, pp. 203–211, Feb. 2021.
- [28] Z. Zhou, M. M. Rahman Siddiquee, N. Tajbakhsh, and J. Liang, "Unet++: A nested U-Net architecture for medical image segmentation," in *Deep Learning in Medical Image Analysis and Multimodal Learning for Clinical Decision Support*. Cham, Switzerland: Springer, 2018, pp. 3–11.
- [29] R. Weteland, K. Engan, T. Eftestøl, V. Kvikstad, and E. A. M. Janssen, "A multiscale approach for whole-slide image segmentation of five tissue classes in urothelial carcinoma slides," *Technol. Cancer Res. Treatment*, vol. 19, Jan. 2020, Art. no. 153303382094678.
- [30] R. A. Khan, Y. Luo, and F.-X. Wu, "RMS-UNet: Residual multi-scale UNet for liver and lesion segmentation," *Artif. Intell. Med.*, vol. 124, Feb. 2022, Art. no. 102231.

- [31] K. Kamnitsas, C. Ledig, V. F. J. Newcombe, J. P. Simpson, A. D. Kane, D. K. Menon, D. Rueckert, and B. Glocker, "Efficient multi-scale 3D CNN with fully connected CRF for accurate brain lesion segmentation," *Med. Image Anal.*, vol. 36, pp. 61–78, Feb. 2017.
- [32] A. Clèrigues, S. Valverde, J. Bernal, J. Freixenet, A. Oliver, and X. Lladó, "Acute ischemic stroke lesion core segmentation in CT perfusion images using fully convolutional neural networks," *Comput. Biol. Med.*, vol. 115, Dec. 2019, Art. no. 103487.
- [33] S. M. Abulnaga and J. Rubin, "Ischemic stroke lesion segmentation in ct perfusion scans using pyramid pooling and focal loss," in *Proc. Int. MICCAI Brainlesion Workshop*. Cham, Switzerland: Springer, 2018, pp. 352–363.
- [34] R. A. Rava, A. R. Podgorsak, M. Waqas, K. V. Snyder, M. Mokin, E. I. Levy, J. M. Davies, A. H. Siddiqui, and C. N. Ionita, "Investigation of convolutional neural networks using multiple computed tomography perfusion maps to identify infarct core in acute ischemic stroke patients," *J. Med. Imag.*, vol. 8, no. 1, Feb. 2021, Art. no. 014505.
- [35] M. Soltanpour, R. Greiner, P. Boulanger, and B. Buck, "Ischemic stroke lesion prediction in CT perfusion scans using multiple parallel U-Nets following by a pixel-level classifier," in *Proc. IEEE 19th Int. Conf. Bioinf. Bioeng. (BIBE)*, Oct. 2019, pp. 957–963.
- [36] M. Soltanpour, R. Greiner, P. Boulanger, and B. Buck, "Improvement of automatic ischemic stroke lesion segmentation in CT perfusion maps using a learned deep neural network," *Comput. Biol. Med.*, vol. 137, Oct. 2021, Art. no. 104849.
- [37] A. S. Kasasbeh, S. Christensen, M. W. Parsons, B. Campbell, G. W. Albers, and M. G. Lansberg, "Artificial neural network computer tomography perfusion prediction of ischemic core," *Stroke*, vol. 50, no. 6, pp. 1578–1581, Jun. 2019.
- [38] L. Tomasetti, S. Hansen, M. Khanmohammadi, K. Engan, L. J. Hølleli, K. D. Kurz, and M. Kampffmeyer, "Self-supervised few-shot learning for ischemic stroke lesion segmentation," 2023, *arXiv:2303.01332*.
- [39] F. Werdiger, M. W. Parsons, M. Visser, C. Levi, N. Spratt, T. Kleinig, L. Lin, and A. Bivard, "Machine learning segmentation of core and penumbra from acute stroke CT perfusion data," *Frontiers Neurol.*, vol. 14, pp. 1–14, Feb. 2023.
- [40] L. de Vries, B. J. Emmer, C. B. L. M. Majoie, H. A. Marquering, and E. Gavves, "PerFU-Net: Baseline infarct estimation from CT perfusion source data for acute ischemic stroke," *Med. Image Anal.*, vol. 85, Apr. 2023, Art. no. 102749.
- [41] J. Bertels, D. Robben, D. Vandermeulen, and P. Suetens, "Contra-lateral information CNN for core lesion segmentation based on native CTP in acute stroke," in *Proc. Int. MICCAI Brainlesion Workshop*. Cham, Switzerland: Springer, 2019, pp. 263–270.
- [42] E. de la Rosa, D. Robben, D. M. Sima, J. S. Kirschke, and B. Menze, "Differentiable deconvolution for improved stroke perfusion analysis," in *Proc. Int. Conf. Med. Image Comput. Comput.-Assist. Intervent.* Cham, Switzerland: Springer, 2020, pp. 593–602.
- [43] K. Amador, M. Wilms, A. Winder, J. Fiehler, and N. Forkert, "Stroke lesion outcome prediction based on 4D CT perfusion data using temporal convolutional networks," in *Proc. 4th Conf. Medical Imag. Deep Learn.* (Proceedings of Machine Learning Research), vol. 143, M. Heinrich, Q. Dou, M. de Bruijne, J. Lellmann, A. Schläfer, and F. Ernst, Eds. PMLR, Jul. 2021, pp. 22–33.
- [44] K. Amador, M. Wilms, A. Winder, J. Fiehler, and N. D. Forkert, "Predicting treatment-specific lesion outcomes in acute ischemic stroke from 4D CT perfusion imaging using spatio-temporal convolutional neural networks," *Med. Image Anal.*, vol. 82, Nov. 2022, Art. no. 102610.
- [45] D. Robben, A. M. M. Boers, H. A. Marquering, L. L. C. M. Langezaal, Y. B. W. E. M. Roos, R. J. van Oostenbrugge, W. H. van Zwam, D. W. J. Dippel, C. B. L. M. Majoie, A. van der Lugt, R. Lemmens, and P. Suetens, "Prediction of final infarct volume from native CT perfusion and treatment parameters using deep learning," *Med. Image Anal.*, vol. 59, Jan. 2020, Art. no. 101589.
- [46] A. Buckner, A. M. Boers, J. C. J. Bot, O. A. Berkhemer, H. F. Lingsma, A. J. Yoo, W. H. van Zwam, R. J. van Oostenbrugge, A. van der Lugt, D. W. J. Dippel, Y. B. W. E. M. Roos, C. B. L. M. Majoie, and H. A. Marquering, "Associations of ischemic lesion volume with functional outcome in patients with acute ischemic stroke: 24-hour versus 1-Week imaging," *Stroke*, vol. 48, no. 5, pp. 1233–1240, May 2017.
- [47] P. Schellinger, R. Bryan, L. Caplan, J. Detre, R. Edelman, C. Jaigobin, C. Kidwell, J. Mohr, M. Sloan, and A. Sorensen, "Evidence-based guideline: The role of diffusion and perfusion MRI for the diagnosis of acute ischemic stroke: Report of the therapeutics and technology assessment subcommittee of the American academy of neurology," *Neurol.*, vol. 75, no. 2, p. 177, 2010.
- [48] M. Goyal, "Challenging the ischemic core concept in acute ischemic stroke imaging," *Stroke*, vol. 51, no. 10, pp. 3147–3155, Oct. 2020.
- [49] C. S. Kidwell, J. L. Saver, J. Mattiello, S. Starkman, F. Vinuela, G. Duckwiler, Y. P. Gobin, R. Jahan, P. Vespa, M. Kalafut, and J. R. Alger, "Thrombolytic reversal of acute human cerebral ischemic injury shown by diffusion/perfusion magnetic resonance imaging," *Ann. Neurol.*, vol. 47, no. 4, pp. 462–469, Apr. 2000.
- [50] M.-A. Labeyrie, G. Turc, A. Hess, P. Hervo, J.-L. Mas, J.-F. Meder, J.-C. Baron, E. Touzé, and C. Oppenheim, "Diffusion lesion reversal after thrombolysis: A MR correlate of early neurological improvement," *Stroke*, vol. 43, no. 11, pp. 2986–2991, Nov. 2012.
- [51] B. D. Murphy, A. J. Fox, D. H. Lee, D. J. Sahlas, S. E. Black, M. J. Hogan, S. B. Coutts, A. M. Demchuk, M. Goyal, R. I. Aviv, S. Symons, I. B. Gulka, V. Beletsky, D. Pelz, V. Hachinski, R. Chan, and T.-Y. Lee, "Identification of penumbra and infarct in acute ischemic stroke using computed tomography Perfusion-Derived blood flow and blood volume measurements," *Stroke*, vol. 37, no. 7, pp. 1771–1777, Jul. 2006.
- [52] L. Tomasetti, L. J. Hølleli, K. Engan, K. D. Kurz, M. W. Kurz, and M. Khanmohammadi, "Machine learning algorithms versus thresholding to segment ischemic regions in patients with acute ischemic stroke," *IEEE J. Biomed. Health Informat.*, vol. 26, no. 2, pp. 660–672, Feb. 2022.
- [53] L. Tomasetti, K. Engan, M. Khanmohammadi, and K. D. Kurz, "CNN based segmentation of infarcted regions in acute cerebral stroke patients from computed tomography perfusion imaging," in *Proc. 11th ACM Int. Conf. Bioinf., Comput. Biol. Health Informat.*, Sep. 2020, pp. 1–8.
- [54] L. Tomasetti, M. Khanmohammadi, K. Engan, L. J. Hølleli, and K. D. Kurz, "Multi-input segmentation of damaged brain in acute ischemic stroke patients using slow fusion with skip connection," in *Proc. Northern Lights Deep Learn. Workshop*, vol. 3, Mar. 2022, pp. 1–9.
- [55] M. Najm, H. Kuang, A. Federico, U. Jogiati, M. Goyal, M. D. Hill, A. Demchuk, B. K. Menon, and W. Qiu, "Automated brain extraction from head CT and CTA images using convex optimization with shape propagation," *Comput. Methods Programs Biomed.*, vol. 176, pp. 1–8, Jul. 2019.
- [56] S. M. Smith, "Fast robust automated brain extraction," *Hum. Brain Mapping*, vol. 17, no. 3, pp. 143–155, Nov. 2002.
- [57] N. Gessert, M. Bengs, M. Schlüter, and A. Schlaefer, "Deep learning with 4D spatio-temporal data representations for OCT-based force estimation," *Med. Image Anal.*, vol. 64, Aug. 2020, Art. no. 101730.
- [58] M. Bengs, N. Gessert, and A. Schlaefer, "4D spatio-temporal convolutional networks for object position estimation in OCT volumes," *Current Directions Biomed. Eng.*, vol. 6, no. 1, pp. 1–10, Sep. 2020.
- [59] A. Myronenko, D. Yang, V. Buch, D. Xu, A. Ihsani, S. Doyle, M. Michalski, N. Tenenholtz, and H. Roth, "4d CNN for semantic segmentation of cardiac volumetric sequences," in *Proc. Int. Workshop Stat. Atlases Comput. Models Heart*. Cham, Switzerland: Springer, 2019, pp. 72–80.
- [60] N. Abraham and N. M. Khan, "A novel focal Tversky loss function with improved attention U-Net for lesion segmentation," in *Proc. IEEE 16th Int. Symp. Biomed. Imag. (ISBI)*, Apr. 2019, pp. 683–687.
- [61] F. Milletari, N. Navab, and S.-A. Ahmadi, "V-Net: Fully convolutional neural networks for volumetric medical image segmentation," in *Proc. 4th Int. Conf. 3D Vis. (3DV)*, Oct. 2016, pp. 565–571.
- [62] S. Bai, J. Zico Kolter, and V. Koltun, "An empirical evaluation of generic convolutional and recurrent networks for sequence modeling," 2018, *arXiv:1803.01271*.
- [63] Y. Gal and Z. Ghahramani, "Dropout as a Bayesian approximation: Representing model uncertainty in deep learning," in *Proc. Int. Conf. Mach. Learn.*, 2016, pp. 1050–1059.
- [64] O. Oktay, J. Schlemper, L. Le Folgoc, M. Lee, M. Heinrich, K. Misawa, K. Mori, S. McDonagh, N. Y. Hammerla, B. Kainz, B. Glocker, and D. Rueckert, "Attention U-Net: Learning where to look for the pancreas," 2018, *arXiv:1804.03999*.
- [65] B. Graham, "Spatially-sparse convolutional neural networks," 2014, *arXiv:1409.6070*.

- [66] S. C. Van De Leemput, M. Meijs, A. Patel, F. J. A. Meijer, B. Van Ginneken, and R. Manniesing, "Multiclass brain tissue segmentation in 4D CT using convolutional neural networks," *IEEE Access*, vol. 7, pp. 51557–51569, 2019.
- [67] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," 2014, *arXiv:1412.6980*.
- [68] T. Birsan and D. Tiba, "One hundred years since the introduction of the set distance by Dimitrie Pompeiu," in *Proc. IFIP Conf. Syst. Modeling Optim.* Cham, Switzerland: Springer, 2005, pp. 35–39.
- [69] R. Advani, H. Naess, and M. W. Kurz, "The golden hour of acute ischemic stroke," *Scandin. J. Trauma, Resuscitation Emergency Med.*, vol. 25, no. 1, pp. 1–5, Dec. 2017.
- [70] A. Meretoja, M. Keshkaran, J. L. Saver, T. Tatlisumak, M. W. Parsons, M. Kaste, S. M. Davis, G. A. Donnan, and L. Churilov, "Stroke thrombolysis: Save a minute, save a day," *Stroke*, vol. 45, no. 4, pp. 1053–1058, 2014.
- [71] A. Meretoja, M. Keshkaran, T. Tatlisumak, G. A. Donnan, and L. Churilov, "Endovascular therapy for ischemic stroke: Save a minute—Save a week," *Neurology*, vol. 88, no. 22, pp. 2123–2127, May 2017.
- [72] S. P. Singh, L. Wang, S. Gupta, H. Goli, P. Padmanabhan, and B. Gulyás, "3D deep learning on medical images: A review," *Sensors*, vol. 20, no. 18, p. 5097, Sep. 2020.
- [73] N. Lakomkin, M. Dhamoon, K. Carroll, I. P. Singh, S. Tuhim, J. Lee, J. T. Fifi, and J. Mocco, "Prevalence of large vessel occlusion in patients presenting with acute ischemic stroke: A 10-year systematic review of the literature," *J. NeuroInterventional Surg.*, vol. 11, no. 3, pp. 241–245, Mar. 2019.
- [74] A. Hakim, S. Christensen, S. Winzeck, M. G. Lansberg, M. W. Parsons, C. Lucas, D. Robben, R. Wiest, M. Reyes, and G. Zaharchuk, "Predicting infarct core from computed tomography perfusion in acute ischemia with machine learning: Lessons from the ISLES challenge," *Stroke*, vol. 52, no. 7, pp. 2328–2337, Jul. 2021.



LIV JORUNN HØLLELI received the cand.med. degree from the Faculty of Medicine, University of Bergen, Norway, in 2012. She is currently pursuing the Ph.D. degree with the Electrical Engineering and Computer Science Department, University of Stavanger (UiS), Norway. She is on the project "In Depth Analysis of Perfusion Computed Tomography (CTP) in Patients With Acute Ischemic Stroke." The project is part of a twin project performed in close cooperation with the Ph.D. student on computer science. She is also a Senior Consultant Radiologist with the Stavanger University Hospital, Norway, holding the European Diploma in Neuroradiology. Her primary research includes neuroradiology, with a focus on medical imaging and AI in ischemic stroke. She is a member of the European Society of Neuroradiology, the Norwegian Society of Neuroradiology, and the Norwegian Medical Association.



KATHINKA DÆHLI KURZ received the cand.med. degree from the Albert-Ludwigs University of Freiburg, Germany, in 2000, the first Ph.D. degree from Germany, in 2004, and the second Ph.D. degree in Dynamic MRI: An important tool in problem-solving of breast abnormalities from the University of Oslo, Norway, in 2011. She is a Senior Consultant Radiologist with the Stavanger University Hospital, Norway, holding the European Diploma in Neuroradiology. She also has a Professor II position with the Electrical Engineering and Computer Science Department, University of Stavanger (UiS), Norway. Her primary research interests include ischemic stroke, neuroradiology, and magnetic resonance imaging. She is a member of the European Society of Neuroradiology, the Norwegian Society of Neuroradiology, and the Norwegian Medical Association.



LUCA TOMASETTI (Member, IEEE) received the B.S. degree in computer science from Università di Bologna, Bologna, Italy, in 2015, and the M.S. degree in computer science from the University of Stavanger (UiS), Stavanger, Norway, in 2019, where he is currently pursuing the Ph.D. degree in artificial intelligence with the Computer Science Department. He is a member of the Biomedical Data Analysis Laboratory (BMDLab). He is part of the project "In Depth Analysis of Perfusion Computed Tomography (CTP) in Patients with Acute Ischemic Stroke" with UiS. The project is part of a twin project performed in close cooperation with the Ph.D. candidate in radiology. His research interests include image processing and machine learning, with an emphasis on ischemic stroke.



KJERSTI ENGAN (Senior Member, IEEE) received the B.E. degree in electrical engineering from the Bergen University College, in 1994, and the M.Sc. and Ph.D. degrees in electrical engineering and information technology from the University of Stavanger (UiS), Norway, in 1996 and 2000, respectively. She is a Professor with the Electrical Engineering and Computer Science Department, UiS, where she is also a Leader of the Biomedical Data Analysis Laboratory (BMDLab). Her research interests include signal and image processing and machine learning, with an emphasis on medical applications and in dictionary learning for sparse signal and image representation. She has a particular interest in AI for newborn survival, stroke detection from CTP imaging, and AI in computational pathology. She was a member of the IEEE Image, Video, and Multidimensional Signal Processing Technical Committee (IVMSP). She has served as an Associate Editor and a Senior Area editor for IEEE SIGNAL PROCESSING LETTERS and an Associate Editor for *SIAM Journal on Imaging Sciences* (SIIMS).



MAHDIEH KHANMOHAMMADI (Member, IEEE) was born in Iran. She received the bachelor's and first master's degree in science and technology in bioelectrical engineering from the University of Tehran, the second master's degree in computational engineering from the University of Halmstad, Sweden, and the Ph.D. degree from the University of Copenhagen, Denmark. Following the completion of her Ph.D. study, she continued her academic journey as a Postdoctoral Researcher with the University of Stavanger.

She is an Associate Professor with the University of Stavanger, has made significant contributions to the field of medical image analysis. She joined the Image Processing Group, Computer Science Department, working on a project in collaboration with the neuroscientists at Aarhus University and the Statistical and Mathematical Science Department, Aalborg University. Her Ph.D. research involved analyzing the effect of behavioral stress on the brain and utilizing microscopic images of rat brains. In collaboration with cardiologists with the Stavanger University Hospital, she focused on analyzing angiographic data of the heart to estimate blood velocity.

...