

BTDNet: a Multi-Modal Approach for Brain Tumor Radiogenomic Classification

Dimitrios Kollias ¹, Karanjot Vendal ¹, Priyanka Gadhavi ¹ and Solomon Russom ¹

¹ School of Electronic Engineering and Computer Science, Queen Mary University of London

* Correspondence: d.kollias@qmul.ac.uk

Abstract: Brain tumors pose significant health challenges worldwide, with glioblastoma being one of the most aggressive forms. Accurate determination of the O6-methylguanine-DNA methyltransferase (MGMT) promoter methylation status is crucial for personalized treatment strategies. However, traditional methods are labor-intensive and time-consuming. This paper proposes a novel multi-modal approach, BTDNet, leveraging multi-parametric MRI scans, including FLAIR, T1w, T1wCE, and T2 3D volumes, to predict MGMT promoter methylation status. BTDNet's main contribution involves addressing two main challenges: the variable volume lengths (i.e., each volume consists of a different number of slices) and the volume-level annotations (i.e., the whole 3D volume is annotated and not the independent slices that it consists of). BTDNet consists of four components: i) the data augmentation one (that performs geometric transformations, convex combinations of data pairs and test-time data augmentation); ii) the 3D analysis one (that performs global analysis through a CNN-RNN); iii) the routing one (that contains a mask layer that handles variable input feature lengths), and iv) the modality fusion one (that effectively enhances data representation, reduces ambiguities and mitigates data scarcity). The proposed method outperforms by at least 3.3% -in terms of F1 score- the state-of-the-art methods in the RSNA-ASNR-MICCAI BraTS 2021 Challenge, offering a promising avenue for enhancing brain tumor diagnosis and treatment.

Keywords: RSNA-ASNR-MICCAI BraTS 2021 Challenge; brain tumor radiogenomic classification; prediction of MGMT promoter methylation status; BTDNet; multimodal approach; routing; mask layer; MixAugment; multi-class focal loss; volume-level annotations; variable length data

1. Introduction

Brain tumors are a complex and heterogeneous group of neoplasms that pose a significant health challenge worldwide. Despite considerable progress in our understanding of their molecular and genetic underpinnings, the diagnosis, prognosis, and treatment of brain tumors remain formidable tasks [1,2]. In recent years, the integration of radiological imaging data with genomic information has emerged as a promising avenue in the field of neuro-oncology. This intersection of radiology and genomics, often referred to as "radiogenomics," has the potential to revolutionize the approach to brain tumor characterization and classification [3].

The inherent complexity and diversity of brain tumors, coupled with the limitations of traditional histopathological methods, have spurred the search for non-invasive, complementary approaches to improve the accuracy of diagnosis and prognosis. Radiogenomic analysis leverages advanced imaging techniques such as magnetic resonance imaging (MRI), computed tomography (CT), and positron emission tomography (PET) to extract a wealth of quantitative and qualitative imaging features [4,5]. These features are then correlated with the genetic and molecular characteristics of brain tumors, offering insights into their underlying biology and clinical behavior.

Glioblastoma, the most aggressive and prevalent primary brain tumor in adults, remains a therapeutic challenge despite advances in neuro-oncology [6]. The genetic and

Citation: Kollias, D.; Vendal, K.; Gadhavi, P.; Russom, S. BTDNet: a Multi-Modal Approach for Brain Tumor Radiogenomic Classification from mpMRI Scans. *Appl. Sci.* **2023**, *13*, 0. <https://doi.org/>

Received:

Revised:

Accepted:

Published:

Copyright: © 2023 by the authors. Submitted to [Appl. Sci.](#) for possible open access publication under the terms and conditions of the Creative Commons Attribution (CC BY) license (<https://creativecommons.org/licenses/by/4.0/>).

epigenetic heterogeneity within glioblastoma underscores the urgent need for precise biomarkers to guide personalized treatment strategies. Among these, the methylation status of the O6-methylguanine-DNA methyltransferase (MGMT) promoter has emerged as a critical determinant of response to temozolomide, the standard chemotherapy for glioblastoma. The predictive power of MGMT promoter methylation status in glioblastoma treatment outcomes underscores the significance of accurate and timely determination [6,7]. Traditionally, the determination of MGMT promoter methylation status has relied on labor-intensive and time-consuming laboratory techniques, such as methylation-specific polymerase chain reaction (MSP) and pyrosequencing.

The emergence of AI-driven methods promises to revolutionize this process by analyzing high-dimensional genomic and epigenomic data. In this work, we utilize multi-parametric MRI (mpMRI) scans, to address the prediction of the MGMT promoter methylation status at the pre-operative baseline MRI scans. The mpMRI scans consist of different types of 3D volumes (which constitute different modalities for the developed methodology): a) Fluid Attenuated Inversion Recovery (FLAIR); b) native (i.e., T1-weighted pre-contrast or in other words T1w); c) post-contrast T1-weighted (T1wCE); and d) T2-weighted (T2) 3D volumes.

The challenges with using such 3D volumes of data for prediction of the MGMT promoter methylation status are two-fold: i) the annotations are at volume-level rather than at slice-level, i.e., there exists one annotation for the whole volume; ii) the volumes have variable lengths, i.e., volumes are 3D signals consisting of series of slices, i.e., 2-D images, and each volume consists of a different number of images. The main contribution of this work is that our proposed method effectively addresses and overcomes these two challenges.

Traditional approaches handle 3D signals using 3D CNN architectures that give one prediction per signal [8,9]; however such architectures are very complex with a large number of parameters and require to have been pre-trained with other large 3D databases. In the healthcare domain and the medical imaging world, large 3D databases -annotated for the purpose of interest- are not easy to find and in most of the cases are not publicly available due to privacy issues, regulatory frameworks and policies. Other traditional approaches make a hypothesis and assign the volume-level label to each slice of the volume and then employ CNN-RNN networks to train with the annotated slices. However, the fact that the whole volume has one label does not mean that each slice in the volume exhibits the MGMT promoter methylation status; it could be the case that only some slices display the MGMT promoter methylation status.

In the case where the volumes have variable lengths, traditional approaches use some ad-hoc strategies; they select a fixed volume length and either remove slices when a larger length is met (thus losing information that could be important for the final decision), or duplicate slices when the volume contains a smaller number of slices (this duplication affects negatively the final decision as the model gets biased towards the repeating data) [10–12]. What is more, this ad-hoc way of selecting the fixed volume length needs to be tuned empirically in every different database.

In this paper we propose BrainTumorDetectionNetwork (BTDNet), a multi-modal approach for prediction of MGMT promoter methylation status. We develop a multi-modal approach as leveraging data from multiple modalities helps the method to: i) have a richer data representation, since different modalities provide complementary information; ii) reduce redundancy and ambiguity in data interpretation; iii) mitigate data scarcity.

BTDNet takes as input an mpMRI scan, i.e., a FLAIR 3D volume, a T1w 3D volume, a T1wCE 3D volume, and a T2 3D volume. BTDNet consists of four components: i) the Data Augmentation; ii) the 3D Analysis; iii) the Routing; and iv) the Modality Fusion. At first, while training, each input volume is passed through the Data Augmentation component. That component applies at first geometric transformations to the input volume and then performs a novel augmentation technique to it. Let us note that, in inference, we apply

Test-Time Data Augmentation [13] to the input volume. The transformed input volume is then processed by the 3D Analysis and Routing components.

The 3D Analysis component consists of a CNN network acting as a feature extractor and is applied to each slice of the volume. The features are then fed to an RNN that captures temporal information within the slices of the same volume. The RNN's features are then passed to the Routing component, consisting of a Mask layer, which dynamically selects specific RNN outputs, followed by a dense layer (i.e., fully connected layer equipped with batch normalization [14] and GELU [15]), whose output features are then fed to the Modality Fusion component. The Mask layer is utilized so as to handle variable input feature lengths (due to variable slices per volume) when training the network.

The 3D Analysis and Routing components are identical and share the same weights for each of the four input 3D volumes (modalities). This means that there exist 4 identical 3D Analysis and Routing components. The Modality Fusion component is fed with the outputs of the Routing component of all four modalities; these outputs are first concatenated and then fed to a dense layer that maps them to the same feature space. Then the output layer follows that performs the final classification. Our method is trained using an extension of the Focal Loss [16] for multi-class classification.

The rest of the sections of this paper are as follows. Section 2 provides a comprehensive review of related works that address the problem of prediction of the MGMT promoter methylation status. Section 3 presents the proposed method, along with its novelties. It further describes the utilized dataset and the pre-processing techniques that we used; finally it presents the performance metrics used to evaluate our approach, as well as implementation details regarding our approach. Section 4 presents the rich experimental study in which we compare our method's performance to that of the state-of-the-art, as well as various ablation studies are presented. Finally, Section 5 provides the conclusion as well as presents future work and extensions of the proposed method.

2. Related Work

In this section we describe the methods that participated in the RSNA-ASNR-MICCAI BraTS 2021 Challenge [17] for classifying the tumor's MGMT promoter methylation status from pre-operative baseline mpMRI data of 2,040 patients. Let us note that all methods: i) tackled the variable volume lengths by either sub-sampling or duplicating slices; ii) tackled the annotation being at volume-level either by 3D CNN networks or by CNN-RNN in which the annotation was propagated to each slice within the volume; iii) developed either multimodal approaches (utilizing all modalities of the mpMRI scan) or unimodal ones (utilizing just the FLAIR modality).

The winning method of the RSNA-ASNR-MICCAI BraTS 2021 Challenge (denoted thereafter as 3D-Resnet10-Trick) [8] employed a 3D CNN model utilizing the Resnet10 [18] architecture with the FLAIR modality from the mpMRI. The model processed slices of size 256×256 . A technique named "Best Central Image Trick", was introduced to construct the 3D input volumes to the model. At first the 'best' slice was selected as the central image of the newly constructed 3D volume (the word 'best' denotes the image that contains the largest brain cutaway view). Then the 20 slices that exist before the 'best' slice and the 20 slices that exist after the 'best' slice are selected (the words 'before' and 'after' refer to the slices within the original FLAIR 3D volume). Finally all these slices are depth concatenated to form the input volume to the 3D-Resnet10.

The runner-up solution of the competition (denoted thereafter as EfficientNet-LSTM-mpMRI) [10] leveraged a CNN-RNN architecture in which EfficientNet-B0 [19] was the selected CNN model and LSTM [20] was the selected RNN model. EfficientNet-B0 was pre-trained and used as a feature extractor, whereas the LSTM was trainable. This approach was multimodal and utilized all four modalities: FLAIR, T1w, T1wCE, and T2. A fixed temporal subsampling was performed to produce ten slices from each MRI modality. These slices were depth concatenated into a 4-channel map, with absent MRI types replaced by zero-filled channels. A 2D convolution converted the 4-channel image into a 3-channel

feature map, aligning with EfficientNet's requirement. Stratified 5-fold cross-validation was employed, and ensemble predictions were derived by averaging outputs from all models.

The third-place solution (denoted thereafter as EfficientNet-Aggr) [11] was a multimodal one comprised of four instances of EfficientNet-B3 model, each dedicated to a specific modality. This approach utilized a stratified split based on patient IDs and classes within the training dataset to derive a validation set. Aggregation of results was performed on a per-patient basis, emphasizing the divergence between maximum predictions and the mean, and subsequently comparing the average and minimum predictions. The prediction associated with the most pronounced difference was retained.

The fourth-placed solution (denoted thereafter as YOLO-EfficientNet) [9] combined Object Detection using YOLOv5 [21] and Classification via 2D and 3D EfficientNet variants. Focusing on the T1wCE type, they employed a systematic seven-slice selection after resampling. The YOLOv5 object detection model, trained on hand-annotated images, identified tumor slices. For 2D classification, the best results were from the T1wCE axials with the EfficientNet B3 model, using diverse augmentations (center cropping and adding noise to 3D data) and techniques like Test-Time Augmentation (TTA) and power ensembling. In the absence of detectable tumors, all slices are globed together and passed on as 3D input. The model was trained with stratified 5-fold.

The fifth-place solution (denoted thereafter as stats-EfficientNet) [12] at first sampled ten images from each modality. The mean of each type was computed to yield four 2D images, which were subsequently depth concatenated into a 4-channel image, which was processed through a 1x1 convolution bottleneck, resulting in a 3-channel feature map. This feature map was fed to EfficientNet, to predict the MGMT value. Given the potential noise in the dataset due to its limited sample size, Taylor Cross Entropy loss was employed.

3. Materials and Method

In this section, at first, we present our proposed method, BTDDNet, detailing all the components that it consists of (Data Augmentation; 3D Analysis; Routing; Modality Fusion) as well as the objective function that was used in its development. Then, we provide details on the utilized in this work dataset, which was used in the RSNA-ASNR-MICCAI BraTS 2021 Challenge. We also present and explain all data pre-processing steps that we followed. Next, we present the performance metric that we utilized to assess the performance of our method. Finally, we present all training implementation details corresponding to the development training and evaluation of our method.

3.1. Proposed Method

At first all input 3D volumes for each modality (FLAIR, T1w, T1wCE and T2) are fed to the Data Augmentation component. While training, this component applies a series of geometric transformations to the inputs and then performs a novel augmentation technique (MixAugment [22]) to them. While testing, this component applies Test-Time Data Augmentation [13] to the inputs.

Then all input 3D transformed volumes for each modality (FLAIR, T1w, T1wCE and T2) are padded with black slices to have lengths t_1, t_2, t_3, t_4 (e.g., each input FLAIR volume consists of t_1 slices, whereas each input T1w volume consists of t_2 slices). Each input 3D volume is first fed to the 3D Analysis component. This component consists of a CNN [23] which performs local (per 2-D image/slice) analysis, extracting features from each slice. Then these features are fed to an RNN [24], which is placed on top of the CNN, so as to capture their temporal dependencies. The CNN and RNN networks perform global (per 3D volume) analysis. The RNN output features are then fed to the Routing component, in which they are concatenated and fed to a Mask layer. This step is essential since we have annotations at volume-level (not at slice-level) and thus we know that all slices (and not just independent slices) may convey important information for the final prediction of the network. This step is also important as it is the Mask layer that dynamically selects RNN outputs taking into account the input length, i.e., the 'true' number of slices of the

currently analyzed volume. The output of the Mask layer is then fed to a dense layer which is equipped with batch normalisation [14] and GELU [15].

Let us note that the above described procedure is followed for each modality. For example, the input 3D FLAIR volume is fed to the 3D Analysis component; the input 3D T1 volume is fed to the 3D Analysis component that is identical and shares the same weights with the corresponding component that the FLAIR modality was fed to. Then the output of the 3D Analysis component of the FLAIR modality is fed to the Routing component; also the output of the 3D Analysis component of the T1 modality is fed to the Routing component which is identical and shares the same weights with the corresponding component of the FLAIR modality.

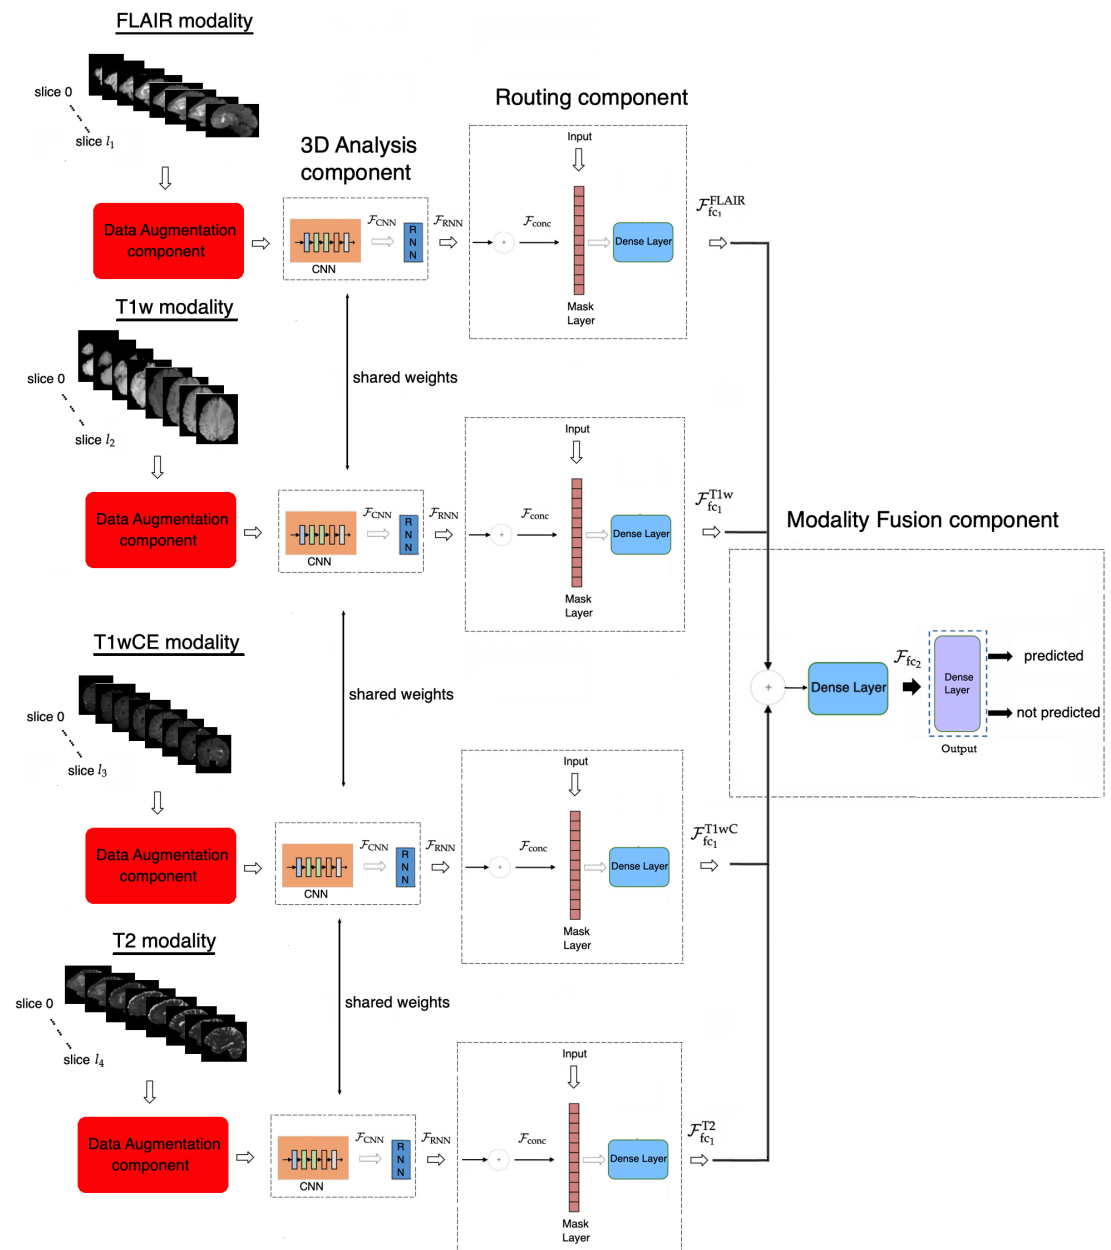


Figure 1. The proposed multimodal BTNet that takes as input mpMRI 3D volumes (FLAIR, T1w, T1wCE and T2 modalities). BTNet consists of the: i) Data Augmentation component (geometric transformations, adapted MixAugment, Test-Time Data Augmentation); ii) 3D Analysis component (CNN plus RNN); iii) Routing component (Mask layer, routing, dense layer); and iv) Modality Fusion component (dense layer, output layer). The symbol '+' in a circle denotes the concatenation operation.

The output features of the Routing component for each modality are fed to the Modality Fusion component. The output features are concatenated and fed to a dense layer that maps them to the same feature space. Finally, the output layer follows providing the final classification. In the following, we further explain in more detail each component of our proposed method. Fig. 1 gives an overview of our proposed framework, BTNet.

3.1.1. Data Augmentation Component

During training, when batches of 3D volumes (X) are sampled (i.e., all slices within the 3D volumes are sampled), at first we perform geometric transformations such as random horizontal flip and rotation and create transformed batches of 3D volumes (T_X). Let us note that we do not perform the same geometric transformation to all slices within the same volume; to the contrary, we apply a different geometric transformation to each slice of the same volume (e.g. one slice is augmented with horizontal flip and 10° rotation, whereas its following slice is augmented only with -4° rotation).

After that, we extend MixAugment [22] to be applied to these transformed batches of 3D volumes. MixAugment is a simple and data-agnostic data augmentation routine that trains a method on convex combinations of pairs of examples and their labels. It extends the training distribution by incorporating the prior knowledge that linear interpolations of feature vectors should lead to linear interpolations of the associated targets. Therefore, it acts as a regularizer to favor linear behavior in-between training examples and it improves the generalization ability of the network, which is always a real world challenge that can hinder the applicability and effectiveness of any method. We extend MixAugment to construct virtual training examples/3D volumes (\tilde{X}, \tilde{y}) (one for each modality) and the method is trained concurrently on both real (r) and virtual (v) examples/3D volumes. The virtual examples for e.g. the FLAIR modality are constructed as follows (similarly are constructed the virtual examples for the other modalities):

$$\begin{aligned}\tilde{X} &= \lambda T_{X_i} + (1 - \lambda) T_{X_j} \\ \tilde{y} &= \lambda y_i + (1 - \lambda) y_j\end{aligned}\quad (1)$$

where: $T_{X_i} \in \mathcal{R}^{H \times W \times t_1}$ and $T_{X_j} \in \mathcal{R}^{H \times W \times t_1}$ are two random FLAIR modality inputs (i.e., 3D volumes) with H and W being the height and width of each slice (i.e., 2D image) and t_1 being the total number of slices -after padding- within the volume; y_i and $y_j \in \{0, 1\}^2$ are their corresponding one-hot labels; $\lambda \sim \text{B}(\alpha, \alpha) \in [0, 1]$ (i.e., Beta distribution) for $\alpha \in (0, \infty)$.

During inference, we apply Test-Time Data Augmentation [13]. For a given input 3D volume (X_i), we create three transformed versions of it: i) one version is when only horizontal flip is performed (F_{X_i}); ii) another version is when random rotation is performed (R_{X_i}); and iii) another version is when both horizontal flip and random rotation is performed (FR_{X_i}). We feed these three versions along with the non-augmented 3D volume to the method which gives a prediction for each version ($p^{F_{X_i}}, p^{R_{X_i}}, p^{FR_{X_i}}, p^{X_i}$, respectively). The final prediction (p_{final}) is the sum of these four outputs ($p_{\text{final}} = p^{F_{X_i}} + p^{R_{X_i}} + p^{FR_{X_i}} + p^{X_i}$). Let us note that we use the logits rather than the probabilities.

3.1.2. 3D Analysis Component

mpMRI 3D volumes consist of 2D slices. CNNs are excellent at capturing local features, patterns, and structures in an image, making them well-suited for extracting meaningful information from individual slices in a 3D MRI volume. By applying 2D convolutions across each slice, the CNN can effectively extract important spatial features from the MRI data. In medical imaging, the spatial context of features is often crucial. MRI scans contain anatomical structures and patterns that may extend across multiple slices. RNNs can capture the spatial dependencies and temporal information within the volume (by analysing the extracted CNN features from each slice). This allows the method to consider

the relationships between slices and how they evolve through the z-axis, which is important for understanding structures or changes/abnormalities over time. This is the aim of the 3D Analysis component, which we describe in more detail next.

After all input 3D volumes for each modality (FLAIR, T1w, T1wCE and T2) have been transformed and then padded to have lengths t_1, t_2, t_3, t_4 , a CNN is applied to each of their slices, performing local (per 2-D image/slice) analysis, extracting features from each slice; e.g. for a slice of a 3D input volume from the FLAIR modality, the extracted features from the CNN are: $\mathcal{F}_{\text{CNN}} \in \mathcal{R}^{H \times W \times D}$, with W, H, D being the feature map's width, height and depth, respectively. Then these features are fed to an RNN, which is placed on top of the CNN so as to capture temporal information and dependencies between consecutive slices within the same volume (for each modality). The RNN analyzes the features of the whole volume sequence, sequentially moving from slice 0 to slice t_1 , if the input volume is for the FLAIR modality (and similarly for the other modalities). As shown in Fig. 1, we get RNN features corresponding to each slice, from 0 to t_1 ; e.g. for a slice of a 3D input volume from the FLAIR modality, the extracted features from the RNN are: $\mathcal{F}_{\text{RNN}} \in \mathcal{R}^V$, with V being the feature vector that is outputted by the RNN; thus for the whole 3D input volume from the FLAIR modality the extracted features from the RNN are: $\mathcal{F}_{\text{RNN}} \in \mathcal{R}^{V \times t_1}$. These RNN features are then fed to the Routing component, described next. In total, the CNN plus RNN perform global (per 3D volume) analysis.

3.1.3. Routing Component

In this component, at first, the RNN features corresponding to the whole input 3D volume are concatenated since our target is the prediction of the MGMT promoter methylation status using the whole volume, similarly to the annotations provided in the utilized dataset. The concatenated features ($\mathcal{F}_{\text{conc}} \in \mathcal{R}^{V \cdot t_1}$) are then fed to the Mask layer. The original (before padding) length l_1 of the input FLAIR volume (similarly l_2, l_3, l_4 for the T1w, T1wCE, T2 volumes) is transferred from the input to the Mask layer to inform the routing process. During model training, the routing mechanism performs dynamic selection of the RNN outputs, selecting as many of them as denoted by the length l of the input volume, to keep their values, while zeroing the values of the rest RNN outputs and thus routing only the selected ones into the following dense layer.

This dense layer is equipped with batch normalization and GELU. We use GELU due to its advantages over ReLU and its variants. The dense layer learns to extract high level information from the concatenated RNN outputs. During training, we update only the weights that connect the dense layer neurons with the RNN outputs routed in the concatenated vector by the Mask layer. The remaining weights are updated whenever (i.e., in another input volume) respective RNN outputs are selected in the concatenated vector by the Mask layer. Objective function minimization is performed, as in networks with dynamic routing, by keeping the weights that do not participate in the routing process constant, and ignoring links that correspond to non-routed RNN outputs. Finally, the output of the dense layer ($\mathcal{F}_{\text{fc}_1} \in \mathcal{R}^{V'}$) is fed to the Modality Fusion component as explained next.

3.1.4. Modality Fusion Component

At first, the output features of the dense layer for the FLAIR modality $\mathcal{F}_{\text{fc}_1}^{\text{FLAIR}} \in \mathcal{R}^{V'}$ are concatenated with the corresponding output features of the dense layer for the other modalities, T1w, T1wCE and T2 ($\mathcal{F}_{\text{fc}_1}^{\text{T1w}} \in \mathcal{R}^{V'}$, $\mathcal{F}_{\text{fc}_1}^{\text{T1wCE}} \in \mathcal{R}^{V'}$, $\mathcal{F}_{\text{fc}_1}^{\text{T2}} \in \mathcal{R}^{V'}$), respectively. These features have already been normalised in the Routing component where batch normalization was performed on the dense layers' outputs for each modality; therefore features coming from different modalities are all within the same range. The concatenated features are then fed to another dense layer ($\mathcal{F}_{\text{fc}_2} \in \mathcal{R}^{V''}$)-equipped with GELU- that maps them to the same feature space. This step is important as leveraging and fusing information from multiple modalities helps the method to have a richer data representation, to reduce redundancy and ambiguity in data interpretation and to mitigate data scarcity. Finally the

output layer follows which consists of two units and gives the final classification for the MGMT promoter methylation status.

3.1.5. Objective Function

For the objective function we built upon the Focal Loss (FL) [16] for multi-class classification, which is defined as follows:

$$\mathcal{L}_{FL} = \sum_{i=1}^{batch} \left[-\alpha(1-p)^\gamma \log p - (1-\alpha)p^\gamma \log(1-p) \right], \quad (2)$$

where: p is the method's estimated probability for the positive class; α balances the importance of positive/negative examples; $(1-p)^\gamma$ and p^γ are modulating factors that reduce the loss contribution from easy examples and extends the range in which an example receives low loss; these modulating factors consist of the tunable focusing parameter $\gamma \geq 0$ that smoothly adjusts the rate at which easy examples are down-weighted.

As explained in the MixAugment strategy above, in each training iteration, the method is fed with both T_{X_i} and T_{X_j} , and the constructed virtual volume \tilde{X} (of Eq. 1). Therefore the overall objective function consists of the sum of the focal losses for the real (r) and virtual (v) volumes:

$$\mathcal{L}_{total} = \mathcal{L}_{FL}^v + \mathcal{L}_{FL}^r + \mathcal{L}_{FL}^j \quad (3)$$

3.2. Dataset

The RSNA-MICCAI consortium has unveiled a dataset [17] that comprises of multi-parametric Magnetic Resonance Imaging (mpMRI) scans (i.e. 3D volumes) from various institutions, annotated for the prediction of a specific genetic characteristic of glioblastoma, namely the MGMT promoter methylation status. Each mpMRI scan consists of four modalities: Fluid attenuated inversion recovery (FLAIR), T1-weighted pre-contrast (T1w), T1-weighted post-contrast (T1wCE), and T2-weighted (T2). Each of these modalities offers a specific imaging perspective, e.g. the FLAIR modality provides imagery post cerebrospinal fluid suppression, where fluidic signals, such as water, are muted to emphasize other components. Fig. 2 shows some slices from each modality of the same mpMRI scan. Each modality contains a varying number of slices. Fig. 3 illustrates the total number of slices of each modality for all mpMRI scans. In total, the dataset comprises of 585 labeled samples, each with dimensions 512×512 .




Figure 2. Example of a whole mpMRI scan, where one can see some slices from: FLAIR modality (top row); T1w modality (second row); T1wCE modality (third row); T2 modality (bottom row).

3.3. Pre-Processing

At first, we performed segmentation of each slice of all mpMRI scans (for all modalities) so as to detect the brain regions and crop them. We used a simple technique for




Figure 3. The total number of slices of each modality for all mpMRI scans.

segmentation that exploited the fact that each slice consists of the brain being in a black background. Each slice is scanned from top to bottom and from left to right and vice-versa in order to find the location (i.e., the coordinates) of the first non-black pixel. We then aggregated all these locations and extracted only four ones that define the created bounding box. We then cropped each 2D slice according to its corresponding bounding box. Additionally, we removed some slices, mainly at the start and end of the MRI scan sequence of slices, since these slices displayed only a minuscule portion of the brain or were entirely void of any useful content. To determine that, we checked whether a slice had a bounding box whose area is less than 5% of the total area of the whole slice; if such cases were found, the corresponding slices were removed. Fig. 4 shows the total number of remaining slices within each mpMRI scan (for each modality). Finally, all segmented slices were resized to $224 \times 224 \times 3$ pixel resolution and their intensity values were normalized to $[-1, 1]$. The resulting segmented slices constitute the input to our method.




Figure 4. The total number of slices, for each modality, within each mpMRI scan.

3.4. Performance Metrics

The performance measure is the average F1 Score across the two categories (i.e., macro F1 Score):

$$\mathcal{P} = \frac{\sum_{expr} F_1^{expr}}{2} \quad (4)$$

The F_1 score [25] is a weighted average of the recall (i.e., the ability of the classifier to find all the positive samples) [26] and precision (i.e., the ability of the classifier not to label as positive a sample that is negative) [25]. The F_1 score takes values in the range $[0, 1]$; high values are desired. The F_1 score is defined as:

$$F_1 = \frac{2 \times \textit{precision} \times \textit{recall}}{\textit{precision} + \textit{recall}} \quad (5)$$

3.5. Implementation Details

We used ResNet18 [18] (discarding the output layer but keeping the global average pooling) as CNN and a single one-directional LSTM [24] consisting of 128 units as RNN. The dense layer in the Routing component consisted of 64 hidden units and the dense layer in the Modality Fusion component consisted of 128 hidden units. We used batch size equal to 4 (we also used values equal to 2, 6, 8, 16, 32 but they all produced worst performance); we used lengths t_1 and t_4 (corresponding to FLAIR and T2 modalities, respectively) equal to 250, lengths t_2 and t_3 (corresponding to T1w and T1wCE modalities, respectively) equal to 200.

We used stratified 5-fold cross validation. The training was divided in two phases: at first, each modality stream/subnetwork was trained independently and then the multi-modal method was trained end-to-end. We utilized SGD with momentum equal to 0.9 [27] and SAM [28] (we also experimented with ADAM optimizer and RMSprop but they both produced worst results); learning rate was 10^{-4} when training from scratch and 10^{-5} when training in an end-to-end manner (we also used learning rate of 10^{-2} , 10^{-3} and 10^{-5} when training from scratch; and we used learning rate of 10^{-3} , 10^{-4} and 10^{-6} when training in an end-to-end manner, but worst performance was produced in all cases). Training was performed on a Tesla V100 32GB GPU.

4. Experimental Results

This section describes a set of experiments evaluating the performance of the proposed approach. At first, we compare the performance of BTDDNet with the performance of the state-of-the-art methods that participated in the RSNA-ASNR-MICCAI BraTS 2021 Challenge; we show that it outperforms these networks by large margins. Finally, we perform ablation studies that illustrate the contribution of the various components of BTDDNet. In particular, we focus on the: 3D Analysis component (i.e., the choice of CNN and RNN models); Routing component (i.e., the existence of the Mask layer, as well as the number of hidden units in the dense layer); contribution of each modality; data augmentations (geometric transformations, MixAugment and Test-Time Data Augmentation); objective function.

4.1. Comparison with the state-of-the-art

At first, we compare the performance of BTDDNet to that of the state-of-the-art methods, the unimodal 3D-Resnet10-Trick [8] (winner of RSNA-ASNR-MICCAI BraTS 2021 Challenge), the multimodal EfficientNet-LSTM-mpMRI [10] (runner up of the Challenge), the multimodal EfficientNet-Aggr [11] (third place in the Challenge), the unimodal YOLO-EfficientNet [9] (fourth place in the Challenge), the multimodal stats-EfficientNet [12] (fifth place in the Challenge). For fair comparison with our method which utilized segmented 3D volumes, we re-implemented these methods and also utilized the same segmented data. Table 1 shows the performance of the state-of-the-art methods, as well as our method's performance. The results are denoted in the form of average of 5 folds \pm spread (i.e., difference between max and min performance within the 5 folds). It can be seen from Table 1 that our method outperforms all other state-of-the-art methods by large margins; in more detail it outperforms the winner of the Challenge, 3D-Resnet10-Trick, by 3.3% and it outperforms all other methods by at least 18.8%. Table 1 also shows that our method achieves the minimum spread among all state-of-the-art methods.

Let us note that all state-of-the-art methods utilized the ad-hoc strategy of selecting fixed input length by removing or duplicating slices within each 3D volume. Let us further note that some methods utilized CNN-RNN architectures in which the 3D volume annotation was propagated to each slice within the volume. Finally, we performed an extra

experiment which we do not show in Table 1 so as to not clutter the results. In more detail, we implemented 3D-Resnet10-Trick (winner of the Challenge) and utilized non-segmented data. Surprisingly we noticed that the method's performance increased from 62.9 ± 4.8 (that it was when segmented data were used) to 63.3 ± 4.4 ; this indicates that the method focused more on regions outside of the brain which played a role in the increase in the performance (however this is not correct as the MGMT promoter methylation status exists in the brain region and should be predicted only from that region).

Table 1. Comparison between BTDDNet and the state-of-the-art on the RSNA-ASNR-MICCAI BraTS 2021 Challenge dataset

Methods	F1 Score
stats-EfficientNet [12]	42.9 ± 6.84
YOLO-EfficientNet [9]	44.4 ± 7.41
EfficientNet-Aggr [11]	46.1 ± 4.68
EfficientNet-LSTM-mpMRI [10]	47.4 ± 4.41
3D-Resnet10-Trick [8]	62.9 ± 4.8
BTDDNet	66.2 ± 3.1

4.2. Ablation Study

In the following, we perform various ablation experiments to evaluate the contribution of each novel component of BTDDNet. At first, we utilize different backbone networks to act as the CNN in the 3D Analysis component of BTDDNet. We experimented with using ResNet50 and ResNet101 (which are bigger networks than ResNet18 which is chosen as the CNN in BTDDNet), EfficientNetB0 and EfficientNetB3 (as has been used by state-of-the-art methods and have proved their value), as well as ConvNeXt-T [29]. As can be seen in Table 2, ResNet18 is a better backbone CNN than its bigger counterparts (ResNet50 and ResNet101); we can see that the bigger the ResNet networks gets (and the more layers it has), the lower the performance is. This occurs probably due to the small size of the training dataset; larger models are too complex to model it. This also occurs in the case of the EfficientNet models as the performance of EfficientNetB0 is higher than that of EfficientNetB3. Table 2 also illustrates that the performance of ConvNeXt-T is higher than that of ResNet50, which was expected. Finally, overall the best performance has been achieved when ResNet18 is utilized. As a second ablation experiment, we utilize different networks to act as the RNN in the 3D Analysis component of BTDDNet. We experimented with using LSTM and GRU, also varying their number of units from 64 to 128 and 256. As can be seen in Table 2, LSTM and GRU produce similar performance. The best results have been obtained when LSTM with 128 units had been used. If the number of units increased, the performance of the method started to decrease.

Next, we assess the contribution of the Routing component and its subcomponents. In more detail, we experimented with not using the Routing component at all or without using the Mask Layer and thus our method was tackling the problem with the different lengths of each volume as a traditional ad-hoc strategy by either removing slices when a larger length is met or duplicating slices when the volume contains a smaller number of slices. Table 2 proves that these approaches do not yield good results; our method outperforms them by large margins (at least 3.3%), as was the case with the state-of-the-art methods that used such ad-hoc strategies. This was expected as removing of slices results in loss of important information for the final decision and duplication of slices negatively affects the final decision and makes the method gets biased towards the repeating data and thus towards one category. We further experimented with using different number of hidden units in the dense layer of the Routing component; we experimented with using 32, 64 and 128 units. The best results have been obtained when we used 64 units.

In the following, we assess the contribution of each modality and thus of the multimodal approach. Table 2 illustrates the results that our method achieves when it is a

Table 2. Ablation Study for the contribution of each component of BTDDNet on the RSNA-ASNR-MICCAI BraTS 2021 Challenge dataset

BTDDNet	F1 Score
ResNet50 as CNN	64.1 ± 3.6
ResNet101 as CNN	62.8 ± 3.88
EfficientNetB0 as CNN	63.9 ± 3.92
EfficientNetB3 as CNN	62.9 ± 4.21
ConvNeXt-T as CNN	64.9 ± 3.7
LSTM, 64 units as RNN	65.1 ± 3.5
LSTM, 256 units as RNN	64.3 ± 4.3
GRU, 64 units as RNN	65 ± 3.7
GRU, 256 units as RNN	64.4 ± 4.4
no Routing component	60.9 ± 5.7
no Mask Layer	62.9 ± 5.2
Dense layer (Routing component), 32 units	65.1 ± 3.4
Dense layer (Routing component), 128 units	64.9 ± 3.6
only FLAIR modality	64.5 ± 3.3
only T1w modality	63.2 ± 4.3
only T1wCE modality	63.8 ± 4.1
only T2 modality	64.1 ± 3.8
no Test-Time Data Augmentation	65.2 ± 3.4
no geometric transformations	65.3 ± 3.3
no MixAugment	64.7 ± 3.8
Categorical Cross Entropy as Objective Function	64.9 ± 3.4
Binary Cross Entropy as Objective Function	64.5 ± 3.5
BTDDNet	66.2 ± 3.1

unimodal approach (using either the FLAIR, or the T1w, or the T1wCE, or the T2 modalities) and when it is a multimodal approach (that utilizes all available modalities). One can see in Table 2, that among the unimodal methods, when the FLAIR modality is used, the best performance is achieved; when T2 modality is used, the second best performance is achieved; when T1wCE modality is achieved, it outperforms the case when T1w is used. Table 2 shows that when the multimodal approach is used, it outperforms all unimodal approaches. This was expected as different modalities provide complementary information. By incorporating multiple modalities, our method could access a richer and more diverse set of data, leading to a more comprehensive understanding of the targeted problem. Additionally, our multimodal method could probably reduce redundancy and ambiguity in data interpretation, becoming more robust and less susceptible to errors or noise. Finally, in some cases, there could be limited data available in a single modality and thus a multimodal approach can help mitigate data scarcity issues by leveraging information from other modalities.

In the following, we perform an ablation study in regards to the proposed data augmentation strategies. In more detail, Table 2 shows the performance of our method when: i) no Test-Time Data Augmentation had been performed (i.e., during inference the method was evaluated only on the real non-augmented data and three augmented versions of the data had not been created and used in the evaluation); ii) no geometric transformations had been used (i.e. both random horizontal flip and rotation had not been performed on the data during training); iii) no MixAugment had been used (i.e., the method was trained only using the real training data and during training no virtual data had been constructed and subsequently used along with the real data). It can be seen that in all cases, the performance of our method decreases, verifying that each data augmentation strategy alone brings a performance gain, and the best performance is achieved when all are used together. The biggest decrease in performance is observed when the MixAugment

strategy is not used. Let us also note that we performed one more experiment which we do not show to not clutter the results; we performed only the geometric transformations (i.e. random horizontal flip and rotation) and the same transformation has been performed on the slices withing the same volume. Our method's performance in this case was 0.8% lower than its corresponding performance when a different random geometric transformation has been performed on the slices withing the same volume.

Finally, the last ablation experiment is with respect to the contribution of the objective function. We experiment with using the proposed multi-class focal loss and the standard categorical cross entropy (when there are two outputs in the method) and binary cross entropy (when there is one output in the method) losses. Table 2 shows that when multi-class focal loss was utilized, our method achieved the best performance, which was at least 1.3% higher than the corresponding one when the other losses were used.

5. Conclusions & Future Work

In this paper we proposed BTDDNet, a new multimodal approach which harmonizes analysis of 3-D image volumes consisting of different number of slices and annotated per volume. In more detail, BTDDNet accepts as input a mpMRI (i.e., a FLAIR 3D volume, a T1w 3D volume, a T1wCE 3D volume and a T2 3D volume) and predicts the MGMT promoter methylation status. BTDDNet consists of three components: the 3D Analysis, the Routing, the Modality Fusion and the Data Augmentation ones.

When BTDDNet is fed with a new 3D volume input, at first a CNN performs local (per 2-D image/slice) analysis, extracting features from each slice. Then these features are fed to an RNN, placed on top of the CNN, so as to capture their temporal dependencies. The CNN and RNN networks perform global (per 3D volume) analysis. The RNN output features are then concatenated and fed to a Mask layer which dynamically selects RNN outputs taking into account their input length. The output of the Mask layer is then fed to a dense layer. This procedure is followed for each of the four modalities; the output features of the dense layer for each modality are then concatenated and fed to a dense layer that maps them to the same feature space. Finally, the output layer follows providing the final classification. The objective function for training BTDDNet was based upon the Focal Loss for multi-class classification.

Excellent performance has been achieved on the dataset of the RSNA-ASNR-MICCAI BraTS 2021 Challenge verifying our developments and surpassing all state-of-the-art methods by large margins. In more detail, BTDDNet outperforms -in terms of F1 score- by 3.3% the winning method of the Challenge and by 18.8%-23.3% the other top-four methods. Finally, an extensive ablation study has been conducted and presented, validating the contribution and performance gain that each component of BTDDNet brings.

Let us mention that in order for our method to be used in real-life clinical situations a much bigger and more diverse dataset will need to be used for training BTDDNet and clinical trials should be conducted.

The hypothesis and limitation of our method is the assumption that all modalities are available and present. An interesting future direction would be to extend BTDDNet to work in the case of one or multiple missing modalities. Another future work could be to develop a joint method that at first performs brain and tumor segmentation and then predicts the MGMT promoter methylation status; this new method can be built upon BTDDNet. Finally, our work only considers the temporal dependencies between same-modality data, while the relationship between different modality data is not investigated. This inter-modality relation is another intriguing future direction.

Author Contributions: Conceptualization and methodology, D. Kollias; software, validation, investigation and formal analysis, D. Kollias, K. Vendal, P. Gadhavi and S. Russom; writing (original draft preparation, review, editing), D. Kollias, K. Vendal, P. Gadhavi and S. Russom; supervision and project administration, D.Kollias. All authors have read and agreed to the published version of the manuscript.

Funding: This research received no external funding

Data Availability Statement: The dataset supporting the conclusions of this article are available at: <https://www.kaggle.com/c/rsna-miccai-brain-tumor-radiogenomic-classification>.

Conflicts of Interest: The authors declare no conflict of interest.

References

1. Shergalis, A.; Bankhead, A.; Luesakul, U.; Muangsin, N.; Neamati, N. Current challenges and opportunities in treating glioblastoma. *Pharmacological reviews* **2018**, *70*, 412–445.
2. McKinnon, C.; Nandhabalan, M.; Murray, S.A.; Plaha, P. Glioblastoma: clinical presentation, diagnosis, and management. *Bmj* **2021**, *374*.
3. Thawani, R.; McLane, M.; Beig, N.; Ghose, S.; Prasanna, P.; Velcheti, V.; Madabhushi, A. Radiomics and radiogenomics in lung cancer: a review for the clinician. *Lung cancer* **2018**, *115*, 34–41.
4. Jansen, R.W.; van Amstel, P.; Martens, R.M.; Kooi, I.E.; Wesseling, P.; de Langen, A.J.; Menke-Van der Houven, C.W.; Jansen, B.H.; Moll, A.C.; Dorsman, J.C.; et al. Non-invasive tumor genotyping using radiogenomic biomarkers, a systematic review and oncology-wide pathway analysis. *Oncotarget* **2018**, *9*, 20134.
5. Villanueva-Meyer, J.E.; Mabray, M.C.; Cha, S. Current clinical brain tumor imaging. *Neurosurgery* **2017**, *81*, 397.
6. Rivera, A.L.; Pelloski, C.E.; Gilbert, M.R.; Colman, H.; De La Cruz, C.; Sulman, E.P.; Bekele, B.N.; Aldape, K.D. MGMT promoter methylation is predictive of response to radiotherapy and prognostic in the absence of adjuvant alkylating chemotherapy for glioblastoma. *Neuro-oncology* **2010**, *12*, 116–121.
7. Louis, D.N.; Perry, A.; Reifenberger, G.; Von Deimling, A.; Figarella-Branger, D.; Cavenee, W.K.; Ohgaki, H.; Wiestler, O.D.; Kleihues, P.; Ellison, D.W. The 2016 World Health Organization classification of tumors of the central nervous system: a summary. *Acta neuropathologica* **2016**, *131*, 803–820.
8. Baba, F. RSNA-MICCAI Brain Tumor Radiogenomic Classification. <https://www.kaggle.com/competitions/rsna-miccai-brain-tumor-radiogenomic-classification/discussion/281347>, 2021.
9. Roberts, D. RSNA-MICCAI Brain Tumor Radiogenomic Classification. <https://www.kaggle.com/competitions/rsna-miccai-brain-tumor-radiogenomic-classification/discussion/280029>, 2021.
10. Phan, M. RSNA-MICCAI Brain Tumor Radiogenomic Classification, 2021.
11. Soares, C. RSNA-MICCAI Brain Tumor Radiogenomic Classification. <https://www.kaggle.com/competitions/rsna-miccai-brain-tumor-radiogenomic-classification/discussion/287713>, 2021.
12. Tangirala, B. RSNA-MICCAI Brain Tumor Radiogenomic Classification. <https://www.kaggle.com/competitions/rsna-miccai-brain-tumor-radiogenomic-classification/discussion/281911>, 2021.
13. Radosavovic, I.; Dollár, P.; Girshick, R.; Gkioxari, G.; He, K. Data distillation: Towards omni-supervised learning. In Proceedings of the Proceedings of the IEEE conference on computer vision and pattern recognition, 2018, pp. 4119–4128.
14. Santurkar, S.; Tsipras, D.; Ilyas, A.; Madry, A. How does batch normalization help optimization? *Advances in neural information processing systems* **2018**, *31*.
15. Hendrycks, D.; Gimpel, K. Gaussian error linear units (gelus). *arXiv preprint arXiv:1606.08415* **2016**.
16. Lin, T.Y.; Goyal, P.; Girshick, R.; He, K.; Dollár, P. Focal loss for dense object detection. In Proceedings of the Proceedings of the IEEE international conference on computer vision, 2017, pp. 2980–2988.
17. Baid, U.; Ghodasara, S.; Mohan, S.; Bilello, M.; Calabrese, E.; Colak, E.; Farahani, K.; Kalpathy-Cramer, J.; Kitamura, F.C.; Pati, S.; et al. The rsna-asnr-miccai brats 2021 benchmark on brain tumor segmentation and radiogenomic classification. *arXiv preprint arXiv:2107.02314* **2021**.
18. He, K.; Zhang, X.; Ren, S.; Sun, J. Deep residual learning for image recognition. In Proceedings of the Proceedings of the IEEE conference on computer vision and pattern recognition, 2016, pp. 770–778.
19. Tan, M.; Le, Q. Efficientnet: Rethinking model scaling for convolutional neural networks. In Proceedings of the International conference on machine learning. PMLR, 2019, pp. 6105–6114.
20. Hochreiter, S.; Schmidhuber, J. Long short-term memory. *Neural computation* **1997**, *9*, 1735–1780.
21. Jocher, G. YOLOv5 by Ultralytics **2020**. <https://doi.org/10.5281/zenodo.3908559>.
22. Psaroudakis, A.; Kollias, D. Mixaugument & mixup: Augmentation methods for facial expression recognition. In Proceedings of the Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, 2022, pp. 2367–2375.
23. LeCun, Y.; Bengio, Y.; et al. Convolutional networks for images, speech, and time series. *The handbook of brain theory and neural networks* **1995**, *3361*, 1995.
24. Gers, F.A.; Schmidhuber, J.; Cummins, F. Learning to forget: Continual prediction with LSTM. *Neural computation* **2000**, *12*, 2451–2471.
25. Powers, D.M. Evaluation: from precision, recall and F-measure to ROC, informedness, markedness and correlation. *arXiv preprint arXiv:2010.16061* **2020**.
26. Schütze, H.; Manning, C.D.; Raghavan, P. *Introduction to information retrieval*; Vol. 39, Cambridge University Press Cambridge, 2008.
27. Amari, S.i. Backpropagation and stochastic gradient descent method. *Neurocomputing* **1993**, *5*, 185–196.

28. Foret, P.; Kleiner, A.; Mobahi, H.; Neyshabur, B. Sharpness-aware minimization for efficiently improving generalization. arXiv preprint arXiv:2010.01412 **2020**.
29. Liu, Z.; Mao, H.; Wu, C.Y.; Feichtenhofer, C.; Darrell, T.; Xie, S. A convnet for the 2020s. In Proceedings of the Proceedings of the IEEE/CVF conference on computer vision and pattern recognition, 2022, pp. 11976–11986.

Disclaimer/Publisher’s Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.