



## OPEN A self-supervised multimodal deep learning approach to differentiate post-radiotherapy progression from pseudoprogression in glioblastoma

Ahmed Gomaa<sup>1,2,3</sup>✉, Yixing Huang<sup>1,2,3</sup>, Pluvio Stephan<sup>1,2,3</sup>, Katharina Breininger<sup>5</sup>, Benjamin Frey<sup>1,2</sup>, Arnd Dörfler<sup>2,4</sup>, Oliver Schnell<sup>2,7</sup>, Daniel Delev<sup>2,7</sup>, Roland Coras<sup>7</sup>, Anna-Jasmina Donaubauer<sup>1,2,3</sup>, Charlotte Schmitter<sup>1,2,3</sup>, Jenny Stritzelberger<sup>8</sup>, Sabine Semrau<sup>1,2,3</sup>, Andreas Maier<sup>9</sup>, Siming Bayer<sup>9</sup>, Stephan Schönecker<sup>3,10</sup>, Dieter H Heiland<sup>6,7,11</sup>, Peter Hau<sup>12,13</sup>, Udo S. Gaipl<sup>1,2,3</sup>, Christoph Bert<sup>1,2,3</sup>, Rainer Fietkau<sup>1,2,3</sup>, Manuel A. Schmidt<sup>2,3,4,14</sup> & Florian Putz<sup>1,2,3,14</sup>

Accurate differentiation of pseudoprogression (PsP) from True Progression (TP) following radiotherapy (RT) in glioblastoma patients is crucial for optimal treatment planning. However, this task remains challenging due to the overlapping imaging characteristics of PsP and TP. This study therefore proposes a multimodal deep-learning approach utilizing complementary information from routine anatomical MR images, clinical parameters, and RT treatment planning information for improved predictive accuracy. The approach utilizes a self-supervised Vision Transformer (ViT) to encode multi-sequence MR brain volumes to effectively capture both global and local context from the high dimensional input. The encoder is trained in a self-supervised upstream task on unlabeled glioma MRI datasets from the open BraTS2021, UPenn-GBM, and UCSF-PDGM datasets (n = 2317 MRI studies) to generate compact, clinically relevant representations from FLAIR and T1 post-contrast sequences. These encoded MR inputs are then integrated with clinical data and RT treatment planning information through guided cross-modal attention, improving progression classification accuracy. This work was developed using two datasets from different centers: the Burdenko Glioblastoma Progression Dataset (n = 59) for training and validation, and the GlioCMV progression dataset from the University Hospital Erlangen (UKER) (n = 20) for testing. The proposed method achieved competitive performance, with an AUC of 75.3%, outperforming the current state-of-the-art data-driven approaches. Importantly, the proposed approach relies solely on readily available anatomical MRI sequences, clinical data, and RT treatment planning information, enhancing its clinical feasibility. The proposed approach addresses the challenge of limited data availability for PsP and TP differentiation and could allow for improved clinical decision-making and optimized treatment plans for glioblastoma patients.

<sup>1</sup>Department of Radiation Oncology, University Hospital Erlangen, Friedrich-Alexander-Universität Erlangen-Nürnberg, Erlangen 91054, Germany. <sup>2</sup>Comprehensive Cancer Center Erlangen-EMN (CCC ER-EMN), Erlangen 91054, Germany. <sup>3</sup>Bavarian Cancer Research Center (BZKF), Erlangen 91052, Germany. <sup>4</sup>Institute of Neuroradiology, University Hospital Erlangen, Friedrich-Alexander-Universität Erlangen-Nürnberg, Erlangen 91054, Germany. <sup>5</sup>Center for Artificial Intelligence and Data Science, Universität Würzburg, Würzburg 97074, Germany. <sup>6</sup>Translational Neurosurgery, Alexander-Friedrich-Universität Erlangen-Nürnberg, Erlangen 91054, Germany. <sup>7</sup>Department of Neurosurgery, University Hospital Erlangen, Friedrich-Alexander-Universität Erlangen-Nürnberg, Erlangen 91054, Germany. <sup>8</sup>Department of Neurology, University Hospital Erlangen, Friedrich-Alexander-Universität Erlangen-Nürnberg, Erlangen 91054, Germany. <sup>9</sup>Pattern Recognition Lab, Friedrich-Alexander-Universität Erlangen-Nürnberg, Erlangen, Germany. <sup>10</sup>Department of Radiation Oncology, University Hospital Ludwig Maximilian University of Munich, 81377 Munich, Germany. <sup>11</sup>Department of Neurological Surgery, Northwestern University Feinberg School of Medicine, Chicago, IL 60611, USA. <sup>12</sup>Department of Neurology, University Hospital Regensburg, Regensburg, Germany. <sup>13</sup>Wilhelm Sander-NeuroOncology Unit, University Hospital Regensburg, Regensburg, Germany. <sup>14</sup>M. A. Schmidt and F. Putz contributed equally to this work. ✉email: ahmed.gomaa@uk-erlangen.de

Glioblastoma is the most frequent primary brain tumor among adults<sup>1</sup>, and understanding its treatment prognosis is critical for optimal treatment selection and patient management<sup>2–5</sup>. Despite advancements in surgical techniques, radiotherapy (RT), and chemotherapy, the prognosis for glioblastoma patients remains poor, with median survival times ranging from 12 to 15 months<sup>6,7</sup>. A particularly immense challenge in the management of glioblastoma is distinguishing between pseudoprogression (PsP) and true progression (TP) following chemoradiation therapy. PsP refers to RT-related transient worsening of radiographic images that mimics tumor growth but does not indicate actual tumor progression and usually resolves spontaneously. TP, in contrast, refers to real tumor growth and proliferation, necessitating a change in a patient's treatment. The treatment options for TP, including surgery and re-irradiation, carry a significant risk of toxicity. Moreover, PsP being mainly induced by RT-related vascular damage, inflammation, and necrosis, can significantly worsen with a second course of radiation<sup>8</sup>.

Misinterpretations can lead to premature discontinuation of effective treatments in the case of PsP or delayed intervention in the case of TP, potentially worsening patient outcomes. Therefore, the precise differentiation between PsP and TP is crucial for optimal clinical decision-making. While surgical biopsy serves as a reliable approach for early tumor progression diagnosis, it is not without limitations. The invasive nature of tissue biopsy can represent a major concern, limiting repeated procedures. Moreover, this concern is further compounded by potential inaccuracies resulting from biopsy site selection, and mixed histological patterns can considerably limit its diagnostic accuracy<sup>9,10</sup>. Furthermore, a biopsy may not be feasible for certain anatomical regions with post-treatment imaging changes<sup>9</sup>. These limitations highlight the need for complementary, non-invasive predictive tools with improved accuracy and reliability.

With the rapid advancement of artificial intelligence, there has been a growing interest in applying machine learning and deep learning methods to improve diagnostic accuracy and treatment planning using medical data<sup>11,12</sup>. For instance, H. Akbari et al. employed a pre-trained Convolutional Neural Network (CNN) to automatically extract features from multiparametric MRI scans. These features were subsequently concatenated with extracted radiomics features to classify progression status using Support Vector Machines (SVM). This approach harnesses the power of deep learning to extract and analyze complex features from imaging data, showing promising results with higher accuracy rates compared to manual interpretations<sup>13</sup>. Using multiparametric MRI data, Lee et al. proposed a hybrid deep learning architecture for differentiating high-grade glioma (HGG) tumor progression from treatment-related changes. Their proposed model consisted of a CNN integrated with Long Short-Term Memory (LSTM) units and was trained on a dataset of 43 biopsy-proven HGG patients. Five standard MRI sequences (DWI, T2, FLAIR, T1 pre- and post-contrast) were combined with two calculated sequences (T1 post-contrast—T1 and T2—FLAIR) to provide additional information for model training. The authors demonstrated that incorporating all available modalities led to the best performance, achieving a mean area under the receiver operating characteristic curve (AUC) of 0.81 in a three-fold cross-validation setup. However, this evaluation did not include an external test set<sup>14</sup>. Moreover, Moassefi et al. leveraged transfer learning by utilizing a pre-trained 3D-DenseNet-121 architecture, originally developed for the sequence registration task. This pre-trained model was subsequently finetuned on a dataset of 124 cases in a five-fold cross-validation. Their approach achieved an AUC of 0.88, though their evaluation was also not conducted on an external test set<sup>15</sup>.

Other works have investigated a multimodal approach that leverages the diagnostic information existing in the clinical data as well as the imaging data. To this end, Sun et al. developed a random forest model to discriminate PsP from TP. Their model incorporated radiomics features extracted from T1 post-contrast MRI scans together with clinical information, such as patient demographics (sex, age), performance status (KPS score), extent of surgical resection, neurological deficits, and mean radiation dose<sup>16</sup>. Additionally, Jang et al. proposed a CNN-LSTM hybrid architecture that was trained using 9 transversal post-contrast T1 image slices, that were centered on the growing lesion, in combination with clinical information consisting of age, gender, molecular features, and RT dose and fractionation parameters<sup>17</sup>.

Although the current literature has shown promising advancements in the progression classification using imaging data, it has relied heavily on first-order Radiomics features and CNNs for processing medical images, which may be suboptimal for capturing long-range dependencies in high-dimensional MR inputs<sup>18</sup>. Additionally, the existing research on combining the imaging modality with clinical data is still limited, with the existing work relying on simple modality fusion techniques. This can lead to limited integration of complementary information across modalities, resulting in suboptimal feature representations and reduced accuracy<sup>17,19</sup>. Furthermore, current methods often address data scarcity through transfer learning, which may not fully leverage the growing number of MRI datasets, as it relies on features from unrelated natural image domains. Self-supervised learning (SSL) for glioblastoma progression classification remains underexplored, despite its success in other medical imaging tasks like disease classification, segmentation, and MRI reconstruction<sup>20–23</sup>. SSL presents an opportunity to develop data-efficient models that learn meaningful features directly from medical imaging data.

Therefore, this work proposes a novel deep learning approach combining a multimodal transformer-based architecture with a self-supervised MRI encoder. Unlike previous approaches that predominantly relied on CNNs, our methodology harnesses the attention mechanism capabilities of transformers to process high-dimensional MRI data more effectively. The attention mechanism of transformers enables them to capture both global and local context information within the data, making them well-suited for tasks involving complex spatial relationships<sup>18,24,25</sup>. Furthermore, the proposed work adopts a self-supervised learning paradigm, which offers several distinct advantages over traditional supervised approaches<sup>26,27</sup>. Self-supervised learning allows the model to learn from a large amount of unlabeled data, which is readily available from various data sources. This is particularly advantageous in the context of medical imaging data and progression classification, where labeled data is scarce and difficult to obtain. We hypothesize that by leveraging unlabeled data, self-

supervised learning could effectively exploit the rich information present in MRI scans, resulting in enhanced generalization and robustness in progression status classification<sup>27</sup>. Additionally, the proposed model employs cross-modal attention to integrate patients' structural clinical features with corresponding imaging data, for enhanced predictive accuracy.

In summary, our work introduces three key contributions to the field of glioblastoma progression classification:

- A multimodal transformer-based architecture for differentiating PsP from TP is proposed and evaluated, demonstrating the advantage of the self- and cross-attention mechanism.
- A self-supervised learning approach is adopted, benefiting from the large amounts of openly available unlabeled MRI data to overcome the data scarcity limitation and improve model generalizability.
- We demonstrate that the proposed deep learning approach combining a multimodal transformer-based architecture with a self-supervised MRI encoder achieves improved predictive performance over the state-of-the-art models as well as simple transfer learning.

To assess model performance and generalizability, we externally validate the proposed transformer architecture on an independent test set from a separate institution.

## Methods

### Datasets

#### *Self-supervised training*

For the pre-text task, three public datasets for glioma cases were used ( $n = 2317$ ). These datasets encompass the University of California San Francisco Preoperative Diffuse glioma MRI (UCSF-PDGM) dataset, comprising 496 cases<sup>28,29</sup>, the Brain Tumor Segmentation (BraTS 2021) dataset with 1251 cases<sup>30–32</sup>, and the University of Pennsylvania Glioblastoma (UPenn-GBM) dataset, which includes 570 cases<sup>29,33</sup>.

#### *Progression classification*

For the training and validation of the proposed deep learning model, we employ the Burdenko's Glioblastoma Progression Dataset<sup>34</sup> ( $n = 180$  patients). The first follow-up scans were obtained approximately one month after the completion of RT, with subsequent scans conducted at roughly three-month intervals. Patients were eligible for inclusion if they were diagnosed with either PsP or TP and had at least two morphologic MR series: a T1-weighted contrast-enhanced (T1CE) and a T2-Fluid-Attenuated Inversion Recovery (FLAIR) sequence, both acquired during the follow-up period when progression was identified. Out of the 180 patients in the dataset, 59 met these criteria. The MR volumes are supplemented with clinical information, including IDH mutation status, O[6]-methylguanine-DNA methyltransferase (MGMT) promoter methylation, age, gender, as well as the dates of RT and follow-up scans. Each patient also has a corresponding RT study, which includes four morphologic MRI sequences (T1, T1CE, T2, and FLAIR), the planning CT scan, and the DICOM RT planning objects, including the three-dimensional dose distribution. Furthermore, to calculate the dose parameters for the enlarging tumor in the follow-up scans, we first segmented the tumors in the follow-up MR volumes using an nnU-Net, trained on the BraTS2021 dataset<sup>35</sup>. Subsequently, we rigidly registered the T1CE and FLAIR follow-up MR sequences at which the tumor enlargement was observed in the radiotherapy planning CT using Advanced Registration Tools (ANTs) version 0.5.3<sup>36</sup>. Leveraging the calculated transform, we mapped the follow-up lesion segmentation into the frame-of-reference of the planning-CT and RT dose distribution to extract the dose-related features, namely mean, min, median, and D98 dose.

A second independent dataset from the prospective Glioblastoma Progression (NCT02600065, Glioblastoma Progression) dataset,  $n = 20$  acquired at the University Hospital Erlangen (UKER) was used as an external test set<sup>37,38</sup>. The Glioblastoma Progression study follows a prospective design with predefined imaging time points obtained through multiparametric MRI. In this study, follow-up scans were conducted every three months after the completion of radiotherapy. This dataset includes similar clinical covariates and imaging modalities, allowing for external validation of the model's performance in a real-world clinical setting. PsP and TP were differentiated in the Glioblastoma Progression dataset either by histology (biopsy or tumor resection,  $n = 8$ ) or by longitudinal imaging follow-up ( $n = 12$ ). The ethics committee at the University Hospital Erlangen approved the conduction of the study (Approval Number: 265\_14 B) and all the patients had given their written informed consent for participation as well as for secondary scientific use of their data. Furthermore, all methods were conducted in accordance with the relevant guidelines and regulations. The patient characteristics of each dataset are included in Table 1. In this supervised learning task, the model was trained using a 5-fold cross-validation approach on Burdenko's dataset, with folds stratified based on progression class. Each fold's trained model was then evaluated on the Glioblastoma Progression dataset. Furthermore, the predictions of the test set are also ensembled using soft voting.

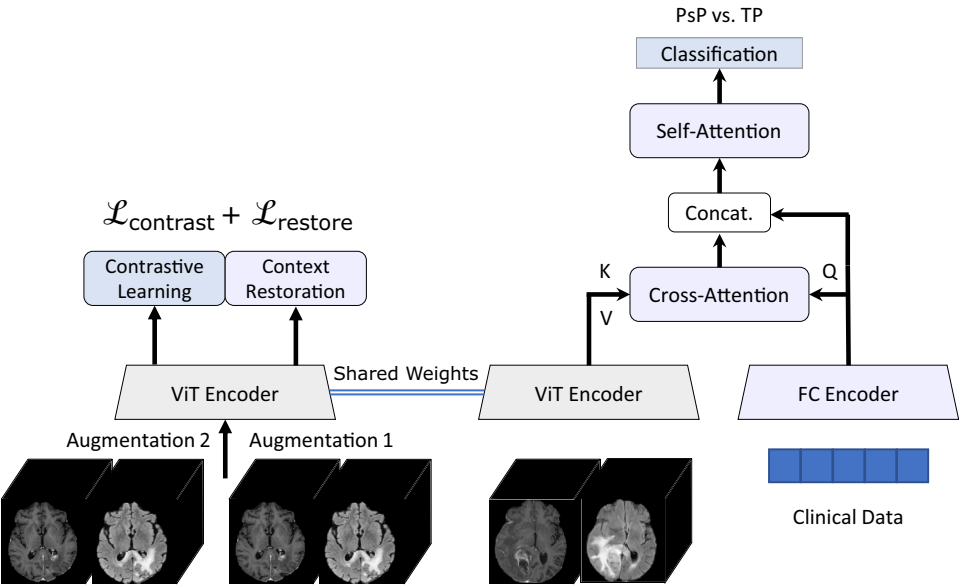
### Data preprocessing

Each subject within the utilized datasets possessed a minimum of 2 corresponding MRI modalities, namely T1 post-contrast and T2-FLAIR sequences. These modalities had already undergone preprocessing steps including spatial normalization, skull stripping, and bias field correction. The two morphological MRI scans underwent further processing before being passed into the deep learning model. This processing involved cropping each scan to a standardized size of  $160 \times 160 \times 160$  voxels while maintaining an isotropic spacing of 1 mm. Subsequently, the cropped scans were concatenated along the channel dimension. Following this, histogram standardization and channel-wise z-normalization were applied to the data.

For the structured clinical features, continuous values were normalized to have a mean of zero and a standard deviation of 1. Categorical values, on the other hand, were encoded to one-hot encoded vectors. Moreover, features that were colinear with different features were dropped.

Variable	Burdenko (n = 59)	GlioCMV UKER (n = 20)
Gender		
Female, n (%)	28 (47.5)	11 (55)
Male, n (%)	31 (52.5)	9 (45)
IDH Mutation Status		
Mutant, n (%)	8 (13.6)	1 (5)
Wildtype, n (%)	36 (61.0)	19 (95)
Unknown, n (%)	15 (25.4)	0 (0)
MGMT Methylation Status		
Methylated, n (%)	14 (23.7)	9 (45)
Unmethylated, n (%)	23 (40.0)	3 (15)
Unknown, n (%)	22 (37.3)	8 (40)
Progression Classification		
True Progression, n (%)	34 (57.6)	11 (55)
Pseudoprogression, n (%)	25 (42.4)	9 (45)
Age in years, median, range	57, (18–82)	58.5, (45–75)

**Table 1.** Patient characteristics of the Burdenko’s Glioblastoma Progression Dataset (Burdenko) and the GlioCMV UKER progression dataset (UKER).



**Fig. 1.** The proposed workflow of the model. On the left-hand side is the self-supervised setup. On the right-hand side is the progression classification setup after the self-supervised training. In the inference phase, multi-parameter MRI volumes and clinical data are fed to the pre-trained ViT encoder and FC encoder, respectively.

To identify the optimum number of features, we performed SHapley Additive exPlanations analysis<sup>39</sup> on the cross-validation sets. After sorting the features based on their importance, as shown in Fig. 3, we chose the first  $M$  features leading to the highest AUC on the validation sets. For this setup, the optimum value of  $M$  was 4.

**Model architecture and training process**

This work introduces a dual-phase deep learning architecture, as depicted in Fig. 1. In the first phase, a Vision Transformer (ViT) serves as an encoder for the MRI data. The parameters of this imaging encoder were optimized in a self-supervised manner using a large collection of unlabeled MRI scans utilizing two proxy tasks, namely context restoration and contrastive learning. This phase aims to employ the self-attention attention capabilities of the ViT encoder in acquiring discriminative and compact representation vectors from complex and high-dimensional MR datasets<sup>24,40</sup>. Detailed information on the self-supervised training strategy is provided in the supplementary material.

In the second phase, the pre-trained self-supervised ViT with frozen parameters is used to encode the input MR volumes. The encoded volumes are represented as  $E_i \in \mathbb{R}^{N \times d}$ , where  $N$  is the number of tokens extracted

from the MR volume, and  $d$  is the dimension on each token. The clinical features are encoded and reshaped into  $E_c \in \mathbb{R}^{M \times d}$ , where  $M$  is the number of tokens derived from the reshaping process and  $d$  is the shared dimension between the clinical and MR tokens. These encoded representations,  $E_i$  and  $E_c$ , are then integrated via guided cross-attention, which captures the pair-wise interactions between the encoded modalities<sup>41–43</sup>. In this process, it is hypothesized that the encoded clinical data guides the cross-attention mechanism by focusing on the most relevant spatial regions in the MR volumes. Specifically, the guided cross-attention mechanism is formulated as follows:

$$\begin{aligned} E_{i \rightarrow c} &= \text{Cross-Attention}_{i \rightarrow c}(E_i, E_c) \\ &= \text{softmax}\left(\frac{Q_c K_i^\top}{\sqrt{d_k}}\right) V_i \\ &= \text{softmax}\left(\frac{W_q E_c (W_k E_i)^\top}{\sqrt{d}}\right) W_v E_i, \end{aligned} \tag{1}$$

where  $\{W_q, W_k, W_v\} \in \mathbb{R}^{d \times d}$  are learnable weight matrices. Subsequently, a self-attention layer is implemented to capture the intra-modal interactions within the vector resulting from the concatenation of  $E_{i \rightarrow c}$  with  $E_c$ . The output of the self-attention module is then fed to two fully connected layers for output prediction after undergoing attention pooling.

Given that in our implementation  $M \ll N$ , where  $M = 4$  and  $N = 512$ , the encoded clinical vector guides the attention process and reduces the token length to  $M$ . This reduction in the token length considerably impacts the model’s computational complexity, particularly when applying the self-attention mechanism in the following stage. Specifically, the complexity of the self-attention process is typically quadratic with respect to the input sequence length. Therefore, by reducing the sequence length from  $N$  to  $M$ , the computation and memory cost is significantly reduced.

Alternative training strategies

We evaluated two commonly used training strategies to compare them with our proposed self-supervised approach. First, we assessed transfer learning, where a Vision Transformer (ViT) was trained on a survival prediction task using 592 samples from the UCSF-PDGM and UPenn-GBM datasets. Afterward, the trained ViT was used as an encoder for the MR data as part of the same multimodal architecture used in the proposed approach in the progression classification task. The second strategy involved training the entire multimodal architecture from scratch end-to-end for the progression classification task.

Results  
Comparison with the state-of-the-art

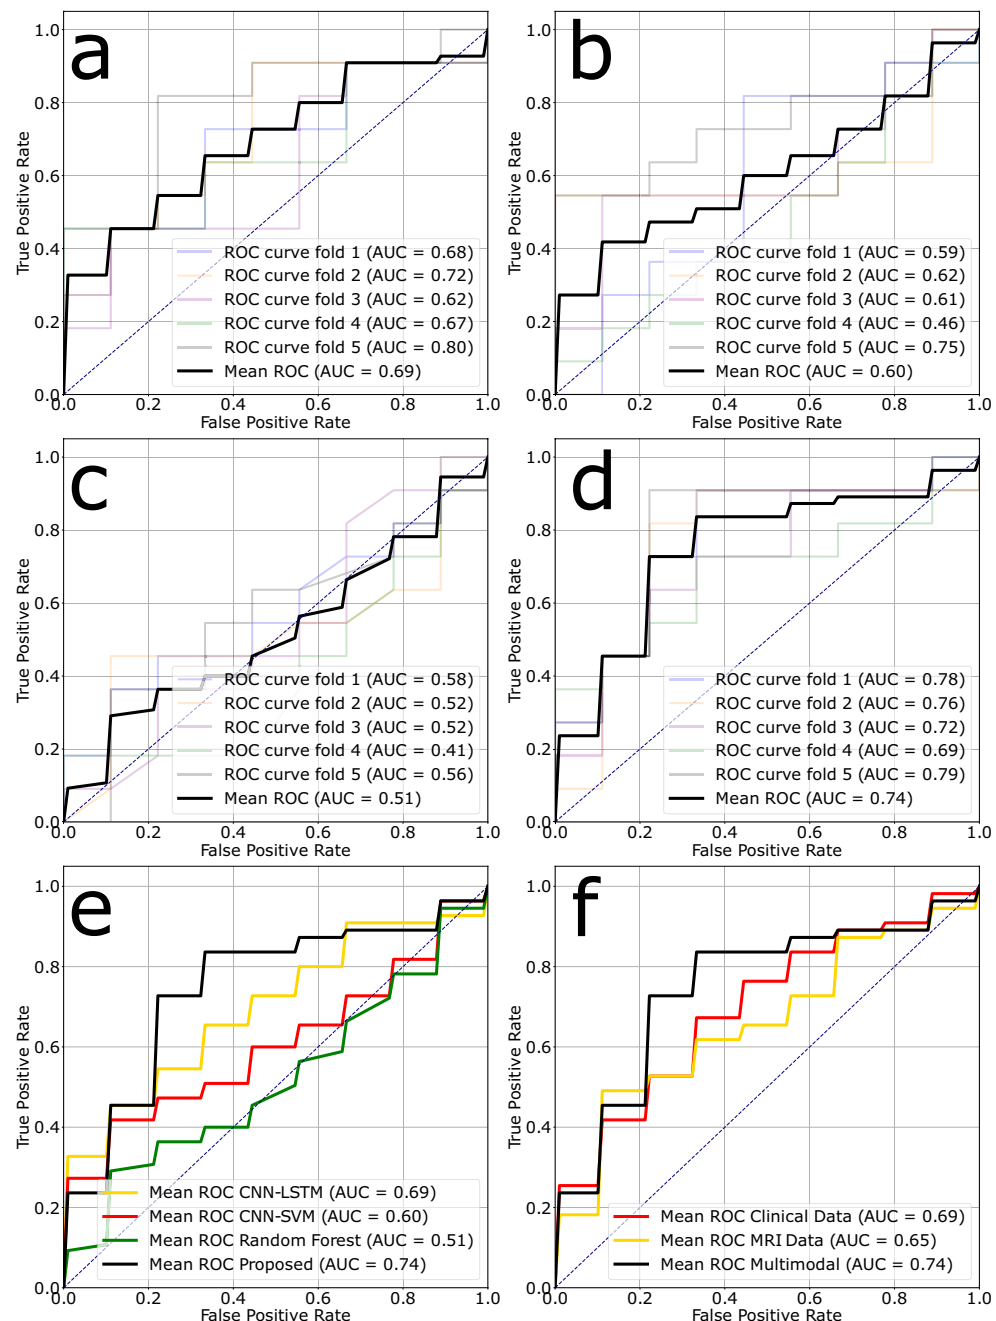
The proposed transformer-based architecture with a pre-trained imaging encoder achieved an AUC of 0.753 on the external test set. This was higher than the CNN-LSTM model by Jang et al. (0.686), the CNN-SVM model by Akbari et al. (0.677), and the Random Forest model by Sun et al. (0.530). In terms of accuracy, the proposed model achieved 0.750, compared to CNN-LSTM at 0.550, CNN-SVM at 0.650, and Random Forest at 0.550. For sensitivity, Random Forest had the highest value at 0.927, but with very low specificity of 0.111. The proposed model had a sensitivity of 0.727, which was higher than the CNN-LSTM model (0.455) and similar to CNN-SVM (0.7). Regarding specificity, the proposed model achieved 0.8, compared to CNN-LSTM at 0.625, CNN-SVM at 0.633, and Random Forest at 0.555. Overall, the model demonstrated an improved AUC compared to the state-of-the-art methods as well as a balanced sensitivity and specificity. These results are summarized in Table 2 and ROC curves are illustrated in Fig. 2(a–e).

Unimodal vs. multimodal

The predictive performance, based on the input modality, was evaluated using AUC, accuracy, sensitivity, and specificity. The multimodal approach outperformed the others, achieving a mean AUC of 0.753, accuracy of

Model	Validation				External test			
	AUC	Accuracy	Sensitivity	Specificity	AUC	Accuracy	Sensitivity	Specificity
CNN-LSTM: Jang et al. <sup>17</sup>	0.744 ± 0.095	0.797 ± 0.104	0.705 ± 0.165	0.880 ± 0.036	0.698 ± 0.060	0.590 ± 0.045	0.491 ± 0.187	0.711 ± 0.241
	–	–	–	–	(0.686)	(0.550)	(0.455)	(0.625)
CNN-SVM: Akbari et al. <sup>13</sup>	0.527 ± 0.199	0.558 ± 0.149	0.619 ± 0.109	0.480 ± 0.271	0.606 ± 0.092	0.570 ± 0.020	0.488 ± 0.093	0.644 ± 0.113
	–	–	–	–	(0.677)	(0.650)	(0.700)	(0.633)
Random Forest: Sun et al. <sup>16</sup>	0.754 ± 0.110	0.730 ± 0.094	0.795 ± 0.144	0.640 ± 0.233	0.516 ± 0.059	0.560 ± 0.021	0.927 ± 0.036	0.111 ± 0.000
	–	–	–	–	(0.530)	(0.550)	(0.909)	(0.555)
Proposed	0.883 ± 0.044	0.770 ± 0.062	0.819 ± 0.109	0.720 ± 0.078	<b>0.748 ± 0.038</b>	<b>0.670 ± 0.090</b>	0.618 ± 0.118	<b>0.733 ± 0.079</b>
	–	–	–	–	(0.753)	(0.750)	(0.727)	(0.800)

**Table 2.** Performance comparison of the proposed model with state-of-the-art models. Results after ensembling with soft voting are provided in parentheses.



**Fig. 2.** ROC curves evaluated on the external test set. Subplot (a) presents the ROC for the CNN-LSTM model (Jang et al.<sup>17</sup>), while (b) displays the ROC for the CNN-SVM model (Akbari et al.<sup>13</sup>). Subplot (c) illustrates the ROC for the random forest model (Sun et al.<sup>16</sup>), and (d) features the proposed model. Subplot (e) highlights the mean ROC curves of the proposed model and existing baselines, and (f) shows the mean ROC curves depending upon the input modality.

0.750, sensitivity of 0.727, and specificity of 0.8. In comparison, the model trained on structured clinical data alone, using a 4-layer fully connected network, reached a mean AUC of 0.727, accuracy of 0.7, sensitivity of 0.636, and specificity of 0.778. When using only imaging data, the finetuned self-supervised ViT model achieved a mean AUC of 0.717, accuracy of 0.7, sensitivity of 0.545, and specificity of 0.857. A detailed summary of these results is presented in Table 3 and Fig. 2(f).

### Impact of self-supervised learning

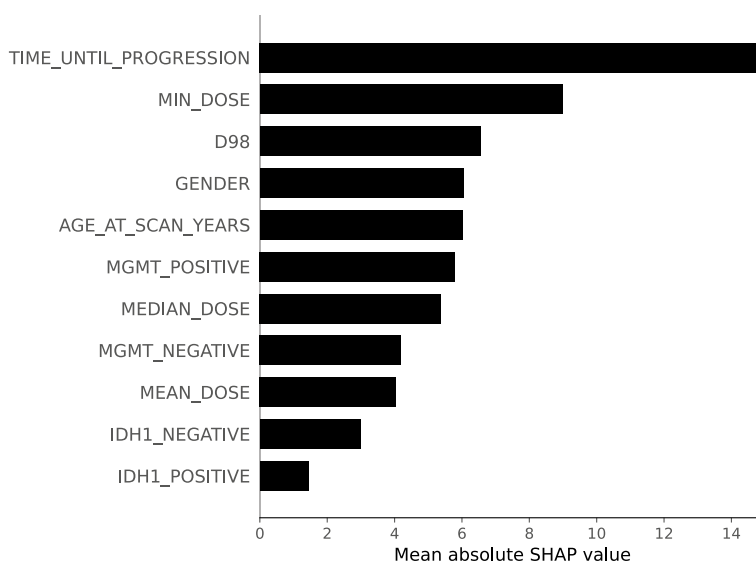
To study the effect of the ViT self-supervised pre-training on the predictive performance, we trained the model under three setups: End-to-end training from scratch, transfer learning, and self-supervised downstream training. For the end-to-end training, the model achieved an AUC of 0.697, an accuracy of 0.6, a sensitivity of

Modality	AUC	Accuracy	Sensitivity	Specificity
MRI	$0.659 \pm 0.065$	$0.620 \pm 0.119$	$0.454 \pm 0.342$	$0.822 \pm 0.160$
	(0.717)	(0.700)	(0.545)	(0.857)
Clinical Data	$0.698 \pm 0.026$	$0.630 \pm 0.072$	$0.618 \pm 0.121$	$0.644 \pm 0.083$
	(0.727)	(0.700)	(0.636)	(0.778)
Multimodal	$0.748 \pm 0.038$	$0.670 \pm 0.090$	$0.618 \pm 0.118$	$0.733 \pm 0.079$
	(0.753)	(0.750)	(0.727)	(0.800)

**Table 3.** Performance on the external test based on input modality. Results after ensembling with soft voting are provided in parentheses.

Model	AUC	Accuracy	Sensitivity	Specificity
ViT End-to-end	$0.675 \pm 0.037$	$0.600 \pm 0.045$	$0.655 \pm 0.134$	$0.533 \pm 0.130$
	(0.697)	(0.600)	(0.545)	(0.667)
ViT Transfer Learning	$0.711 \pm 0.060$	$0.612 \pm 0.037$	$0.582 \pm 0.136$	$0.645 \pm 0.129$
	(0.727)	(0.650)	(0.636)	(0.700)
ViT Self-supervised	$0.748 \pm 0.038$	$0.670 \pm 0.090$	$0.618 \pm 0.118$	$0.733 \pm 0.079$
	(0.753)	(0.750)	(0.636)	(0.800)

**Table 4.** Predictive performance of the multimodal model on the external test set based on the training method of the ViT imaging encoder, with the soft voting ensembling results reported in parentheses.



**Fig. 3.** Mean feature importance analysis using SHAP in a 5-fold cross-validation setup.

0.545, and a specificity of 0.667. In the transfer learning approach, the model was first pre-trained on survival prediction using 592 MR samples and then fine-tuned for the progression classification task. The fine-tuned model achieved an AUC of 0.727, an accuracy of 0.650, a sensitivity of 0.636, and a specificity of 0.7. Lastly, in the self-supervised downstream training setup, the model achieved the highest overall performance, with an AUC of 0.753, an accuracy of 0.75, a sensitivity of 0.636, and a specificity of 0.8. These results, summarized in Table 4, indicate that self-supervised training led to better predictive performance, particularly in terms of AUC and sensitivity, compared to both end-to-end and transfer learning approaches.

### Feature importance analysis

To understand the factors contributing most to the model's predictions, we conducted a feature importance analysis using SHAP analysis on the validation sets. The results illustrated in Fig. 3 show that the Isocitrate dehydrogenase (IDH) gene mutation status has the least impact on the model's outcome. On the other hand, the time between the last treatment and the first progression has the greatest influence on the model's prediction. Furthermore, out of the included dose-related features, the minimum dose, and the near-minimum dose (D98) of the enlarging lesion are the most relevant for progression status classification.

## Discussion

The differentiation of pseudoprogression from true progression is essential for optimal clinical decision-making and patient management following chemoradiation for glioblastoma. The miss-classification of one over the other can result in unnecessary and potentially unfavorable alteration of therapeutic regimes. This involves starting new interventions that may be unnecessary or even hazardous due to side effects, or the premature termination of successful treatments. Moreover, distinguishing between these two events is essential for assessing the effectiveness of therapies in clinical trials, which impacts the precise evaluation of innovative glioblastoma therapy approaches. Novel deep learning techniques could address this critical challenge.

In this work, we present a multimodal transformer-based deep learning model for progression classification in glioblastoma patients. Equipped with the self-supervised learning approach, the model can leverage the more abundant unlabeled brain MR scans to improve classification performance in data-scarce applications. Furthermore, the transformer's self-attention mechanism allows the model to learn the long-term dependencies that inherently exist within high-dimensional inputs, such as MRI data. Moreover, the proposed model utilizes guided cross-modal attention, allowing the model to benefit from the incorporation of different modalities. The proposed model achieves promising predictive performance with an AUC of 0.753, accuracy of 0.75, sensitivity of 0.727, and specificity of 0.8, outperforming the existing state-of-the-art methods (Table 2). The improvement in predictive performance can be attributed to a multitude of factors. The proposed model uses the whole-brain MRI volume to make predictions, compared to a limited number of slices used in the other baselines. This allows the model to benefit from a more comprehensive spatial understanding of the tumor's characteristics. Furthermore, the proposed model employs guided cross-attention to integrate the information from the different input modalities, allowing the model to learn more refined discriminative features, in contrast to the direct concatenation approach used by the baseline models<sup>13,16,17</sup>.

This work investigated the predictive capability of each modality and the impact of integrating them into a unified model. As shown in Table 3, the clinical data outperformed the imaging modality across all metrics except for specificity. Moreover, the results also show that the model has successfully learned to integrate complementary features that are not available when using a single modality, outperforming the models of single modalities.

To investigate the impact of the learning strategy of the MRI encoder on the predictive performance, we compare the performance of the proposed approach with transfer learning as well as end-to-end training. As shown in Table 4, the transfer learning approach performed 0.727, 0.65, 0.636, and 0.7 in terms of AUC, accuracy, sensitivity, and specificity, respectively. The transfer learning method outperformed the training from scratch approach, which achieved 0.697, 0.6, 0.545, and 0.667 for the AUC, accuracy, sensitivity, and specificity, respectively. The notable improvement in the performance provided by the self-supervised encoder model relative to the other methods can be attributed to the ability of the former to learn nuanced features and patterns from the unlabeled data in the pre-training process. This contrasts with the other approaches, which rely on labeled data and may struggle with overfitting or insufficient representation learning, especially when labeled data is limited<sup>26,27</sup>.

The previous results indicate that the model benefits from incorporating the structured clinical features. However, in this context, not all features contribute equally to the deep learning progression classification, which is illustrated in Fig. 3. This graph shows that the time between the treatment termination and the time point at which progression was observed has the greatest impact on the model's output compared to the other clinical features. This observation is aligned with the findings in the literature which states the PsP typically occurs within the first 3 months after the end of the treatment for roughly 60% of progression cases<sup>9,44</sup>. Moreover, among the dose-related features, the minimum and the near-minimum dose (D98) of the enlarging lesion play a considerable role in determining the model's output. However, the other dose-related features have a lesser impact on the prediction outcome. The minimum dose and the D98 (near-minimum dose) of the enlarging lesion can indicate, whether the enlarging lesion is entirely located in the high-dose region of the previous radiation therapy. As the radiotherapy dose causes pseudoprogression but impedes tumor progression, a higher D98 and minimum dose is expected for pseudoprogression than for real tumor progression. For the molecular-pathologic markers, MGMT methylation status has a greater impact on the model's prediction than IDH1. This is also in line with the available literature showing that pseudoprogression is more likely to occur in MGMT-methylated glioblastomas<sup>45</sup>. Taken together, the feature importance analysis indicates that the model is able to successfully learn the most relevant clinical and dosimetric features for differentiating PsP from TP.

While this work provides insights into the task of progression classification, it is essential to mention the limitations to prevent overgeneralization of the reported results. First, as typical for pseudoprogression datasets, the sample size of the two datasets used in this study is relatively limited ( $n = 59$  and  $n = 20$ ), which might introduce biases and limit the generalizability of the findings. Second, the model utilizes only two MRI sequences, namely T1CE and FLAIR. The model could benefit greatly by including additional advanced MRI modalities, such as diffusion-weighted imaging (DWI) and perfusion MRI (PWI), metabolic MRI-CEST<sup>46</sup>, and/or magnetic resonance spectroscopy (MRS), which could provide further insights into tissue characteristics and improve the accuracy and robustness of the model's predictions<sup>47–53</sup>. The third limitation is the relative lack of straightforward interpretability offered by deep learning models.

In conclusion, this work introduces a multimodal transformer-based deep learning model that improves the differentiation between pseudoprogression and true progression in glioblastoma patients in a non-invasive manner. Owing to the self-supervised pre-training, the model can achieve a competitive performance despite the limited training data. Furthermore, the guided cross-attention employed by this model demonstrates an effective approach to integrating information from imaging data, clinical, molecular-pathologic, and dosimetric parameters, resulting in improved predictive accuracy and outperforming the state-of-the-art methods. The results of the proposed work demonstrate the advantage of using readily available anatomical MRI sequences,

clinical data, and RT treatment planning information for progression classification. These findings suggest that the model has the potential to support clinical decision-making by helping to reduce misdiagnosis and guide more effective treatment decisions in real-world settings. Future work could explore the potential performance improvements by incorporating advanced MRI modalities, such as DWI, PWI, and MRS, to further enhance the model's predictive accuracy and robustness in differentiating pseudoprogression from true progression.

## Data availability

The datasets generated during and/or analyzed during the current study are available from the corresponding author upon reasonable request.

Received: 20 December 2024; Accepted: 8 May 2025

Published online: 17 May 2025

## References

- Price, M. et al. CBTRUS statistical report: American brain tumor association & NCI neuro-oncology branch adolescent and young adult primary brain and other central nervous system tumors diagnosed in the united states in 2016–2020. *Neuro-oncology* **26**, iii1–iii53 (2024).
- Ostrom, Q. T. et al. CBTRUS statistical report: Primary brain and other central nervous system tumors diagnosed in the united states in 2013–2017. *Neuro-oncology* **22**, iv1–iv96 (2020).
- Hagag, A. et al. Deep learning for cancer prognosis prediction using portrait photos by stylegan embedding. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, 198–208 (Springer, 2024).
- Zhou, Q. et al. Integrated MRI radiomics, tumor microenvironment, and clinical risk factors for improving survival prediction in patients with glioblastomas. *Strahlentherapie und Onkologie* 1–13 (2024).
- Lau, D., Magill, S. & Aghi, M. Molecularly targeted therapies for recurrent glioblastoma: Current and future targets. *Neurosurg. Focus* **37**(6), E15. <https://doi.org/10.3171/2014.9.FOCUS14519> (2014).
- Stupp, R. et al. Radiotherapy plus concomitant and adjuvant temozolomide for glioblastoma. *N. Engl. J. Med.* **352**, 987–996 (2005).
- Koshy, M. et al. Improved survival time trends for glioblastoma using the seer 17 population-based registries. *J. Neuro-oncol.* **107**, 207–212 (2012).
- Brandtsma, D. & van den Bent, M. J. Pseudoprogression and pseudoresponse in the treatment of gliomas. *Curr. Opin. Neurol.* **22**, 633–638 (2009).
- Da Cruz, L. H., Rodriguez, I., Domingues, R., Gasparetto, E. & Sorensen, A. Pseudoprogression and pseudoresponse: imaging challenges in the assessment of posttreatment glioma. *Am. J. Neuroradiol.* **32**, 1978–1985 (2011).
- Yang, Y. et al. Adding DSC PWI and DWI to BT-RADS can help identify postoperative recurrence in patients with high-grade gliomas. *J. Neuro-Oncol.* **146**, 363–371 (2020).
- Huang, Y. et al. Principles of artificial intelligence in radiooncology. *Strahlentherapie und Onkologie* 1–26 (2024).
- Erdur, A. C. et al. Deep learning for autosegmentation for radiotherapy treatment planning: State-of-the-art and novel perspectives. *Strahlentherapie und Onkologie* 1–19 (2024).
- Akbari, H. et al. Histopathology-validated machine learning radiographic biomarker for noninvasive discrimination between true progression and pseudo-progression in glioblastoma. *Cancer* **126**, 2625–2636 (2020).
- Lee, J. et al. Discriminating pseudoprogression and true progression in diffuse infiltrating glioma using multi-parametric MRI data through deep learning. *Sci. Rep.* **10**, 20331 (2020).
- Moassefi, M. et al. A deep learning model for discriminating true progression from pseudoprogression in glioblastoma patients. *J. Neuro-oncol.* **159**, 447–455 (2022).
- Sun, Y.-Z. et al. Differentiation of pseudoprogression from true progression in glioblastoma patients after standard treatment: A machine learning strategy combined radiomics features from T1-weighted contrast-enhanced imaging. *BMC Med. Imaging* **21**, 1–12 (2021).
- Jang, B.-S., Jeon, S. H., Kim, I. H. & Kim, I. A. Prediction of pseudoprogression versus progression using machine learning algorithm in glioblastoma. *Sci. Rep.* **8**, 12516 (2018).
- Dosovitskiy, A. et al. An image is worth 16x16 words: Transformers for image recognition at scale. arXiv preprint [arXiv:2010.11929](https://arxiv.org/abs/2010.11929) (2020).
- Wang, W., Tran, D. & Feiszli, M. What makes training multi-modal classification networks hard? In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 12695–12705 (2020).
- Huang, S.-C. et al. Self-supervised learning for medical image classification: A systematic review and implementation guidelines. *NPJ Dig. Med.* **6**, 74 (2023).
- Fedorov, A. et al. On self-supervised multimodal representation learning: an application to alzheimer's disease. In *2021 IEEE 18th International Symposium on Biomedical Imaging (ISBI)*, 1548–1552 (IEEE, 2021).
- Qiu, S. et al. Physics-guided self-supervised learning for retrospective t1 and t2 mapping from conventional weighted brain MRI: Technical developments and initial validation in glioblastoma. *Magn. Reson. Med.* **92**, 2683–2695 (2024).
- Pani, K. & Chawla, I. A hybrid approach for multi modal brain tumor segmentation using two phase transfer learning, SSL and a hybrid 3dunet. *Comput. Electr. Eng.* **118**, 109418 (2024).
- He, K. et al. Transformers in medical image analysis. *Intell. Med.* **3**, 59–78 (2023).
- Raghu, M., Unterthiner, T., Kornblith, S., Zhang, C. & Dosovitskiy, A. Do vision transformers see like convolutional neural networks?. *Adv. Neural Inf. Process. Syst.* **34**, 12116–12128 (2021).
- Chen, L. et al. Self-supervised learning for medical image analysis using image context restoration. *Med. Image Anal.* **58**, 101539 (2019).
- Shurrah, S. & Duwairi, R. Self-supervised learning methods and applications in medical imaging analysis: A survey. *PeerJ Comput. Sci.* **8**, e1045 (2022).
- Calabrese, E. et al. The University of California San Francisco preoperative diffuse glioma MRI dataset. *Radiol.: Artif. Intell.* **4**, e220058 (2022).
- Clark, K. et al. The cancer imaging archive (TCIA): Maintaining and operating a public information repository. *J. Dig. Imaging* **26**, 1045–1057 (2013).
- Baid, U. et al. The RSNA-ASNR-MICCAI BraTS 2021 benchmark on brain tumor segmentation and radiogenomic classification (2021). [arXiv:2107.02314](https://arxiv.org/abs/2107.02314).
- Menze, B. H. et al. The multimodal brain tumor image segmentation benchmark (BRATS). *IEEE Trans. Med. Imaging* **34**, 1993–2024 (2014).
- Bakas, S. et al. Advancing the cancer genome atlas glioma MRI collections with expert segmentation labels and radiomic features. *Sci. Data* **4**, 1–13 (2017).
- Bakas, S. et al. The university of pennsylvania glioblastoma (UPenn-GBM) cohort: Advanced MRI, clinical, genomics, & radiomics. *Sci. Data* **9**, 453 (2022).

34. Zolotova, S. V. et al. Burdenko's glioblastoma progression dataset (Burdenko-GBM-Progression) (version 1) (2023). Accessed 10 Apr 2023.
35. Isensee, F., Jaeger, P. F., Kohl, S. A., Petersen, J. & Maier-Hein, K. H. nnU-Net: A self-configuring method for deep learning-based biomedical image segmentation. *Nat. Methods* **18**, 203–211 (2021).
36. Avants, B. B. et al. A reproducible evaluation of ANTs similarity metric performance in brain image registration. *Neuroimage* **54**, 2033–2044 (2011).
37. Goerig, N. L. et al. Early mortality of brain cancer patients and its connection to cytomegalovirus reactivation during radiochemotherapy. *Clin. Cancer Res.* **26**, 3259–3270 (2020).
38. Goerig, N. L. et al. Frequent occurrence of therapeutically reversible CMV-associated encephalopathy during radiotherapy of the brain. *Neuro-oncology* **18**, 1664–1672 (2016).
39. Lundberg, S. M. & Lee, S.-I. A unified approach to interpreting model predictions. *Advances in neural information processing systems* **30** (2017).
40. Azad, R. et al. Advances in medical image analysis with vision transformers: A comprehensive review. *Medical Image Anal.* **91**, 103000 (2023).
41. Lu, J., Batra, D., Parikh, D. & Lee, S. Vilbert: Pretraining task-agnostic visiolinguistic representations for vision-and-language tasks. *Advances in neural information processing systems* **32** (2019).
42. Goma, A. et al. Comprehensive multimodal deep learning survival prediction enabled by a transformer architecture: A multicenter study in glioblastoma. *Neuro-Oncol. Adv.* **6**, vdae122 (2024).
43. Tsai, Y.-H. H. et al. Multimodal transformer for unaligned multimodal language sequences. In *Proceedings of the conference. Association for computational linguistics. Meeting*, Vol. 2019, 6558 (NIH Public Access, 2019).
44. Brandsma, D., Stalpers, L., Taal, W., Sminia, P. & van den Bent, M. J. Clinical features, mechanisms, and management of pseudoprogression in malignant gliomas. *Lancet Oncol.* **9**, 453–461 (2008).
45. Brandes, A. A. et al. MGMT promoter methylation status can predict the incidence and outcome of pseudoprogression after concomitant radiochemotherapy in newly diagnosed glioblastoma patients. *J. Clin. Oncol.* **26**, 2192–2197 (2008).
46. Zhou, J. et al. Review and consensus recommendations on clinical apt-weighted imaging approaches at 3T: Application to brain tumors. *Magn. Reson. Med.* **88**, 546–574 (2022).
47. Ungan, G. et al. Early pseudoprogression and progression lesions in glioblastoma patients are both metabolically heterogeneous. *NMR Biomed.* **37**, e5095 (2024).
48. Chawla, S. et al. Metabolic and physiologic magnetic resonance imaging in distinguishing true progression from pseudoprogression in patients with glioblastoma. *NMR Biomed.* **35**, e4719 (2022).
49. El-Abtah, M. E. et al. Magnetic resonance spectroscopy outperforms perfusion in distinguishing between pseudoprogression and disease progression in patients with glioblastoma. *Neuro-Oncol. Adv.* **4**, vdac128 (2022).
50. Sidibe, I., Tensaouti, F., Roques, M., Cohen-Jonathan-Moyal, E. & Laprie, A. Pseudoprogression in glioblastoma: Role of metabolic and functional MRI-systematic review. *Biomedicines* **10**, 285 (2022).
51. Wen, P. Y. et al. Updated response assessment criteria for high-grade gliomas: Response assessment in neuro-oncology working group. *J. Clin. Oncol.* **28**, 1963–1972 (2010).
52. Henson, J. W., Ulmer, S. & Harris, G. Brain tumor imaging in clinical trials. *Am. J. Neuroradiol.* **29**, 419–424 (2008).
53. Smith, J. S. et al. Serial diffusion-weighted magnetic resonance imaging in cases of glioma: Distinguishing tumor recurrence from postresection injury. *J. Neurosurg.* **103**, 428–438 (2005).

## Author contributions

The study was designed and conceived by FP, YH, and AG. AG wrote the program code and conducted the experimental work. AS contributed to the code. AG, AS, and FP wrote the original draft. SS, FP, AG, CS, MAS, AJD, and BF collected and processed the data. MAS and AD provided the imaging data for this study. MAS and AD provided imaging expertise. CB, KB, RF, AD, OS, RC, DD, JS, AM, SB, SAS, DH, PH, UG, and MAS reviewed and edited the script. All authors have read and approved the script.

## Funding

Open Access funding enabled and organized by Projekt DEAL.

## Declarations

## Competing interests

The authors declare no competing interests.

## Additional information

**Supplementary Information** The online version contains supplementary material available at <https://doi.org/10.1038/s41598-025-02026-7>.

**Correspondence** and requests for materials should be addressed to A.G.

**Reprints and permissions information** is available at [www.nature.com/reprints](http://www.nature.com/reprints).

**Publisher's note** Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

**Open Access** This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.