Spatiotemporal Glioma Progression Prediction Using Longitudinal MRI and Clinical Data

Machine Learning and Pattern Recognition

Team: Gaganpreet, T H S Eswar Reddy, Karri Divya Naidu
Instructor: Dr Siddharth

Problem Statement

Clinical Challenge

Gliomas are the most aggressive primary brain tumors. Even after surgery, chemotherapy, and radiation, they frequently recur. The key challenge is not just that they recur, but predicting when and how aggressively, because post-treatment MRI scans are extremely hard to interpret. Tumor regrowth and treatment side effects look nearly identical, even to experienced radiologists.

Core Question

Given longitudinal post-treatment MRI scans and clinical data, can we model the patient-specific risk of tumour progression over time?

Early Decision-Making

Flag high-risk patients for aggressive follow-up instead of waiting

Personalized Treatment

Tailor therapy based on predicted tumor evolution

Clinical Impact

Impact: shifting glioma management from reactive to proactive — earlier detection, better treatment decisions, and ultimately better patient outcomes.

Literature Review

Paper Title	Added prognostic value of 3D deep learning-derived features from preoperative MRI for adult-type diffuse gliomas
Dataset	5 distinguished datasets
Methodology	• 3D Convolutional Neural Network • MRI images used as input • Deep learning used to learn spatial features automatically
Key Findings	• Deep learning models can learn complex spatial imaging features directly from MRI. • The model improved survival prediction accuracy in glioma patients.
Key Limitations	• Model analyzes single MRI scans. • Does not model temporal tumor evolution across follow-up scans

Lee, J. O., Ahn, S. S., Choi, K. S., Lee, J., Jang, J., Park, J. H., Hwang, I., Park, C. K., Park, S. H., Chung, J. W., & Choi, S. H. (2024). Added prognostic value of 3D deep learning-derived features from preoperative MRI for adult-type diffuse gliomas. Neuro-oncology, 26(3), 571–580. https://doi.org/10.1093/neuonc/noad202

Literature Review - Brain Tumor Classification(2025)

Paper Title	Context aware machine learning techniques for brain tumor classification and detection – A review
Dataset	BraTS dataset
Methodology	• SVM, CNN, and Inception-v3 tested for classification • Features included tumor shape, texture, and intensity statistics • Pearson correlation used to eliminate highly correlated features
Key Findings	• Inception-v3 outperformed SVM and traditional CNN • Imaging features successfully captured tumor heterogeneity • Random Forest used to rank feature importance
Key Limitations	• Each MRI scan treated as independent (no temporal modeling) • No longitudinal data used (single timepoint analysis only)

Amjad, U., Raza, A., Fahad, M., Farid, D., Akhunzada, A., Abubakar, M., & Beenish, H. (2025). Context aware machine learning techniques for brain tumor classification and detection – A review. Heliyon, 11(2), e41835. https://doi.org/10.1016/j.heliyon.2025.e41835

Literature Review -MRI image preprocessing methods

Paper Title	A deep learning framework integrating MRI image preprocessing methods for brain tumor segmentation and classification
Dataset	Four MRI modalities (T1, T1CE, T2, FLAIR) with tumor segmentation masks
Methodology	• UNet for tumor segmentation distinguishing necrosis, edema, and enhancing core • VGG and GoogleNet for tumor grading • Gamma correction and WSO module for preprocessing • SMOTE for class imbalance • Grad-CAM for visualization of influential brain regions
Key Findings	• Multi-modal MRI input significantly improved grading accuracy • Dice coefficient, Precision, Recall used as metrics • Grad-CAM revealed clinically meaningful regions driving predictions
Key Limitations	• Focused on grading only (not longitudinal progression prediction) • No temporal modeling across follow-up scans • Single timepoint per patient

Rehman, A., Naz, S., Razzak, I., Akram, F., & Imran, M. (2022). A deep learning-based framework for automatic brain tumors classification using transfer learning. Smart Health, 26, 100379. https://doi.org/10.1016/j.smhl.2022.100379

Literature Review - Benchmarking of Deep Learning Models

Paper Title	Timepoint-Specific Benchmarking of Deep Learning Models for Glioblastoma Follow-Up MRI
Dataset	Glioblastoma follow-up MRI with multiple post-treatment timepoints
Methodology	• Benchmarked 11 DL families (CNNs, LSTMs, hybrids, Transformers, Mamba state-space models) • Unified QC-driven pipeline with patient-level cross-validation • Data imbalance handled via oversampling and augmentation • 2D Mamba+CNN hybrid as final best model
Key Findings	• 2D Mamba+CNN consistently offered best accuracy-efficiency tradeoff • Second follow-up timepoint gave clearer progression separation than first • CNN-LSTM with multiparametric MRI + clinical data boosted AUC significantly
Key Limitations	• Used only 2 timepoints (did not model full longitudinal sequence) • 2D slices used instead of full 3D volumetric analysis • Clinical metadata not deeply integrated into the model

Guo, W., & Mirzaei, G. (2025). Timepoint-specific benchmarking of deep learning models for glioblastoma follow-up MRI. arXiv. https://doi.org/10.48550/arXiv.2511.18595

Literature Review: Glioma ML Research Summary

Paper	Methodology	Dataset	Key Findings	Limitations
Added prognostic value of 3D deep learning features from preoperative MRI for diffuse gliomas — Lee et al. (2024)	3D CNN; MRI images as input; deep learning for spatial feature extraction	5 distinguished datasets	DL models learn complex spatial features; improved survival prediction accuracy	Single MRI scan only; no temporal tumor evolution modeling
Context aware ML techniques for brain tumor classification and detection — Amjad et al. (2025)	SVM, CNN, Inception-v3; tumor shape/texture/intensity features; Pearson correlation; Random Forest for feature ranking	BraTS dataset	Inception-v3 outperformed SVM and CNN; imaging features captured tumor heterogeneity	No temporal modeling; single timepoint analysis only
Deep learning framework integrating MRI preprocessing for brain tumor segmentation and classification — Rehman et al. (2022)	UNet (segmentation); VGG & GoogleNet (grading); Gamma correction; SMOTE; Grad-CAM visualization	T1, T1CE, T2, FLAIR with segmentation masks	Multi-modal MRI improved grading accuracy; Grad-CAM revealed clinically meaningful regions	Grading only; no longitudinal progression prediction; single timepoint per patient
Timepoint-Specific Benchmarking of Deep Learning Models for Glioblastoma Follow-Up MRI — Guo & Mirzaei (2025)	Benchmarked 11 DL families (CNNs, LSTMs, Transformers, Mamba); QC-driven pipeline; 2D Mamba+CNN hybrid as best model	Glioblastoma follow-up MRI (multiple post-treatment timepoints)	2D Mamba+CNN best accuracy-efficiency tradeoff; CNN-LSTM with multiparametric MRI + clinical data boosted AUC	Only 2 timepoints; 2D slices not 3D volumetric; clinical metadata not deeply integrated

Dataset Overview

MU-Glioma-Post

203

Patients

Glioma patients with longitudinal follow-up

596

Timepoints

Post-treatment MRI scans across multiple visits

74

Clinical Features

Demographics, genomics, treatment outcomes

4

MRI Modalities

T1, T1c, T2, FLAIR sequences per timepoint

```
MU-Glioma-Post/
    PatientID_XXXX/
      Timepoint_X/
        brain_t1n*.nii.gz    ← T1 native
        brain_t1c*.nii.gz    ← T1 contrast-enhanced
        brain_t2w*.nii.gz    ← T2-weighted
        brain_t2f*.nii.gz    ← FLAIR
        tumorMask.nii        ← segmentation mask (labels 1-4)
```

Dataset Overview

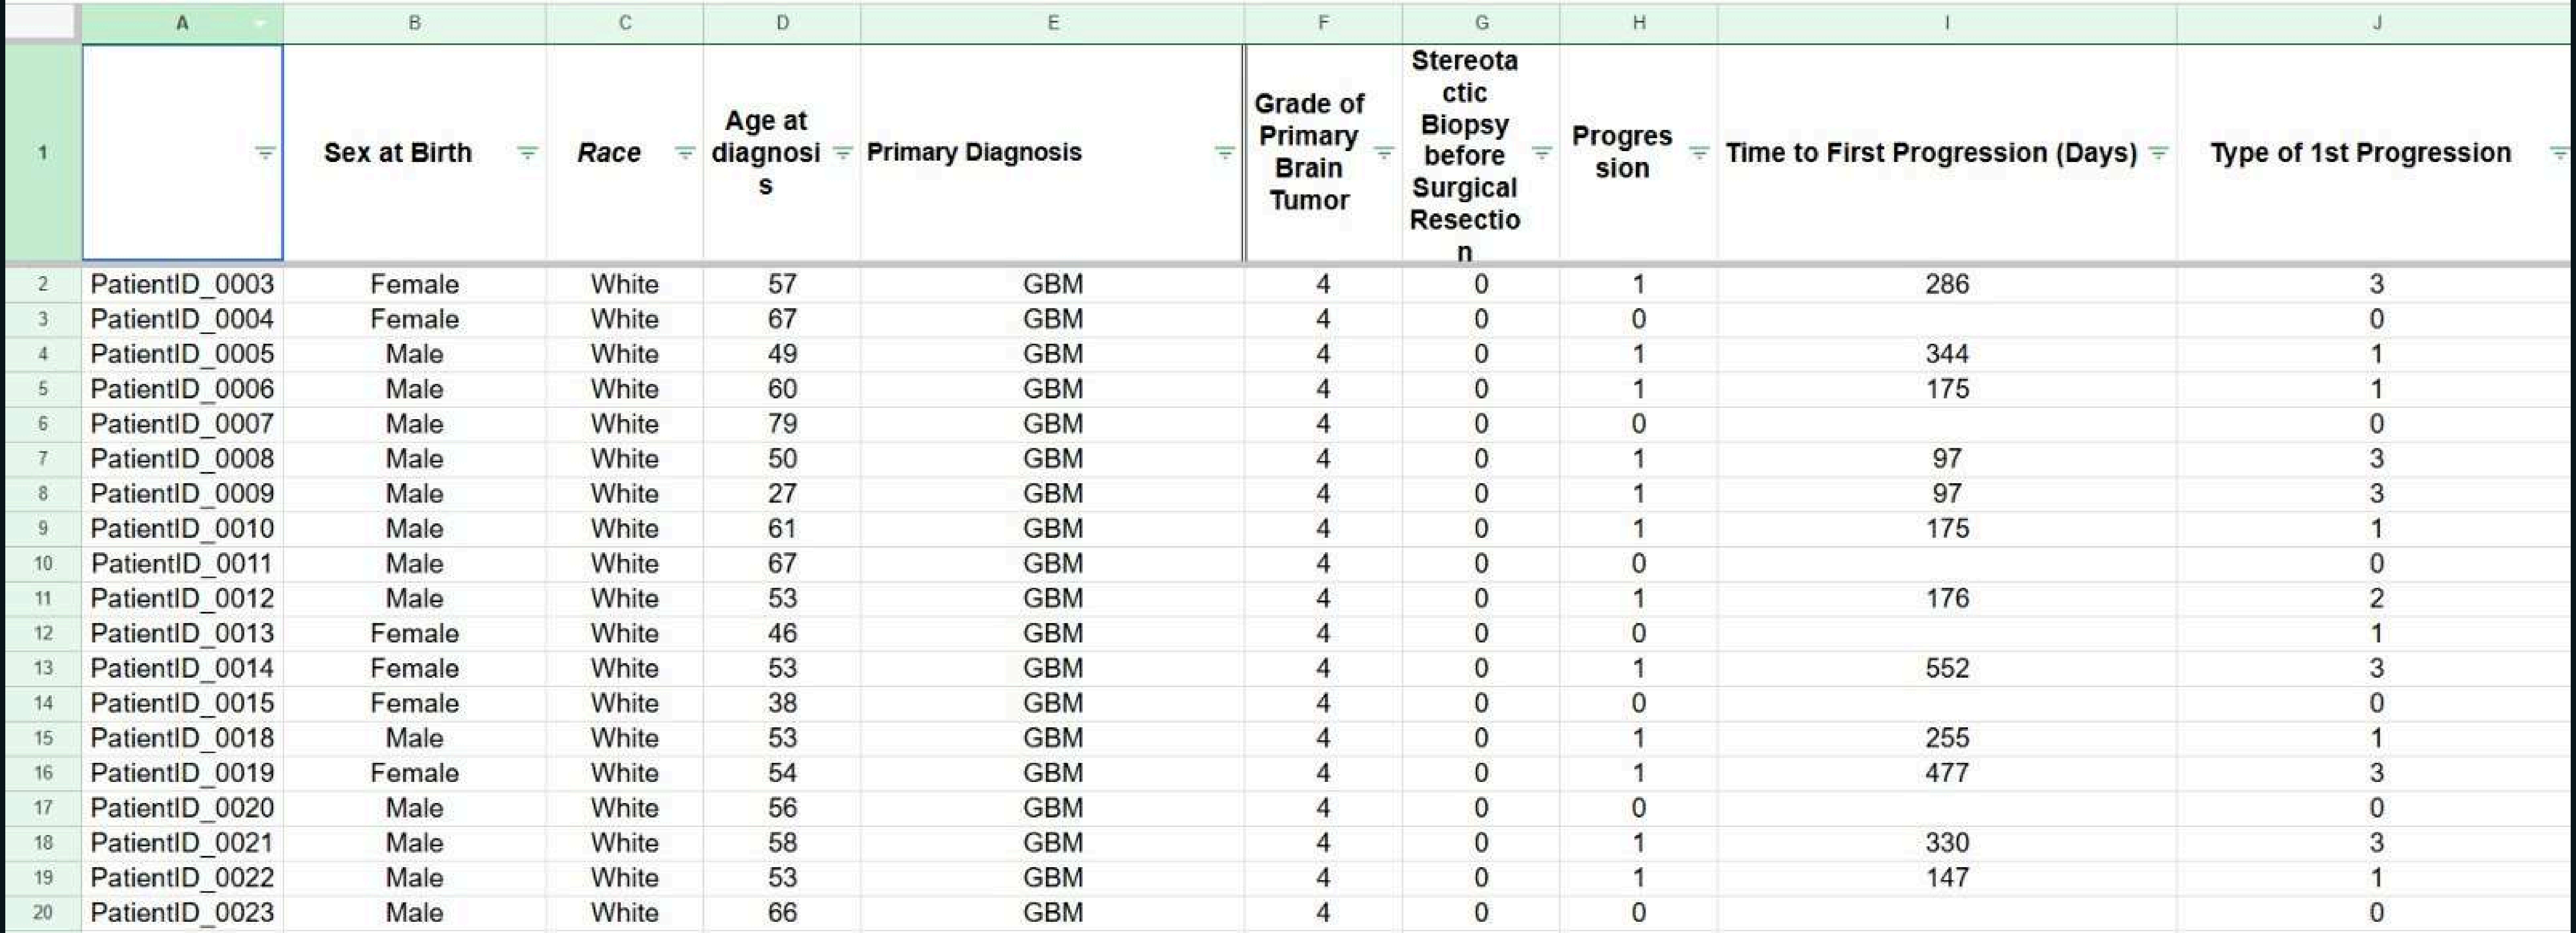

	Sex at Birth	*Race*	Age at diagnosis	Primary Diagnosis	Grade of Primary Brain Tumor	Stereotactic Biopsy before Surgical Resection	Progression	Time to First Progression (Days)	Type of 1st Progression
PatientID_0003	Female	White	57	GBM	4	0	1	286	3
PatientID_0004	Female	White	67	GBM	4	0	0		0
PatientID_0005	Male	White	49	GBM	4	0	1	344	1
PatientID_0006	Male	White	60	GBM	4	0	1	175	1
PatientID_0007	Male	White	79	GBM	4	0	0		0
PatientID_0008	Male	White	50	GBM	4	0	1	97	3
PatientID_0009	Male	White	27	GBM	4	0	1	97	3
PatientID_0010	Male	White	61	GBM	4	0	1	175	1
PatientID_0011	Male	White	67	GBM	4	0	0		0
PatientID_0012	Male	White	53	GBM	4	0	1	176	2
PatientID_0013	Female	White	46	GBM	4	0	0		1
PatientID_0014	Female	White	53	GBM	4	0	1	552	3
PatientID_0015	Female	White	38	GBM	4	0	0		0
PatientID_0018	Male	White	53	GBM	4	0	1	255	1
PatientID_0019	Female	White	54	GBM	4	0	1	477	3
PatientID_0020	Male	White	56	GBM	4	0	0		0
PatientID_0021	Male	White	58	GBM	4	0	1	330	3
PatientID_0022	Male	White	53	GBM	4	0	1	147	1
PatientID_0023	Male	White	66	GBM	4	0	0		0

Dataset Overview

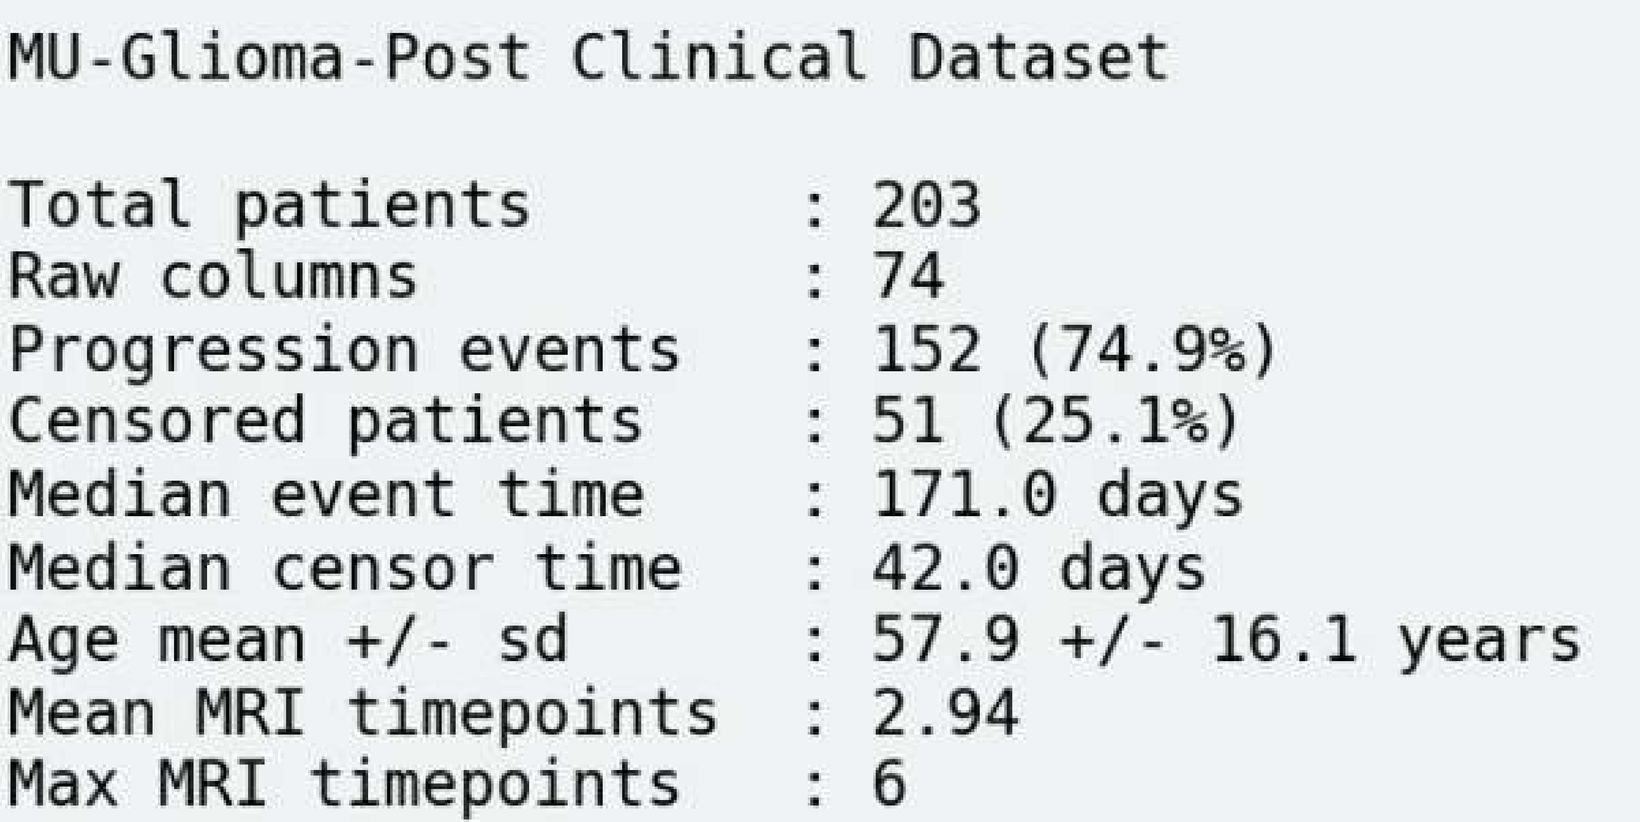

```
MU-Glioma-Post Clinical Dataset

Total patients          : 203
Raw columns             : 74
Progression events      : 152 (74.9%)
Censored patients       : 51 (25.1%)
Median event time       : 171.0 days
Median censor time      : 42.0 days
Age mean +/- sd         : 57.9 +/- 16.1 years
Mean MRI timepoints     : 2.94
Max MRI timepoints      : 6
```

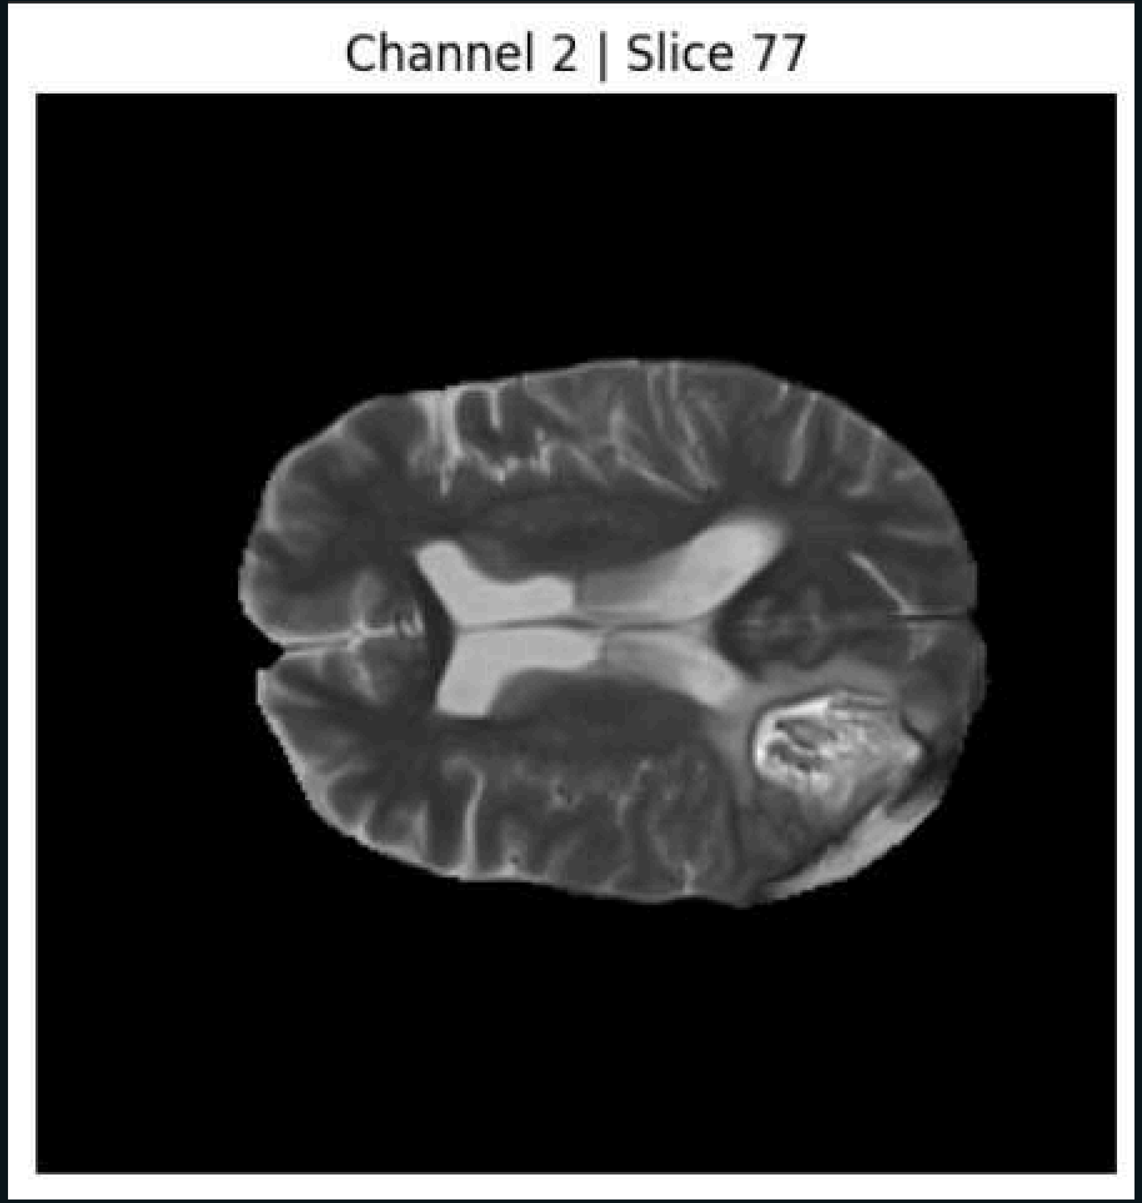

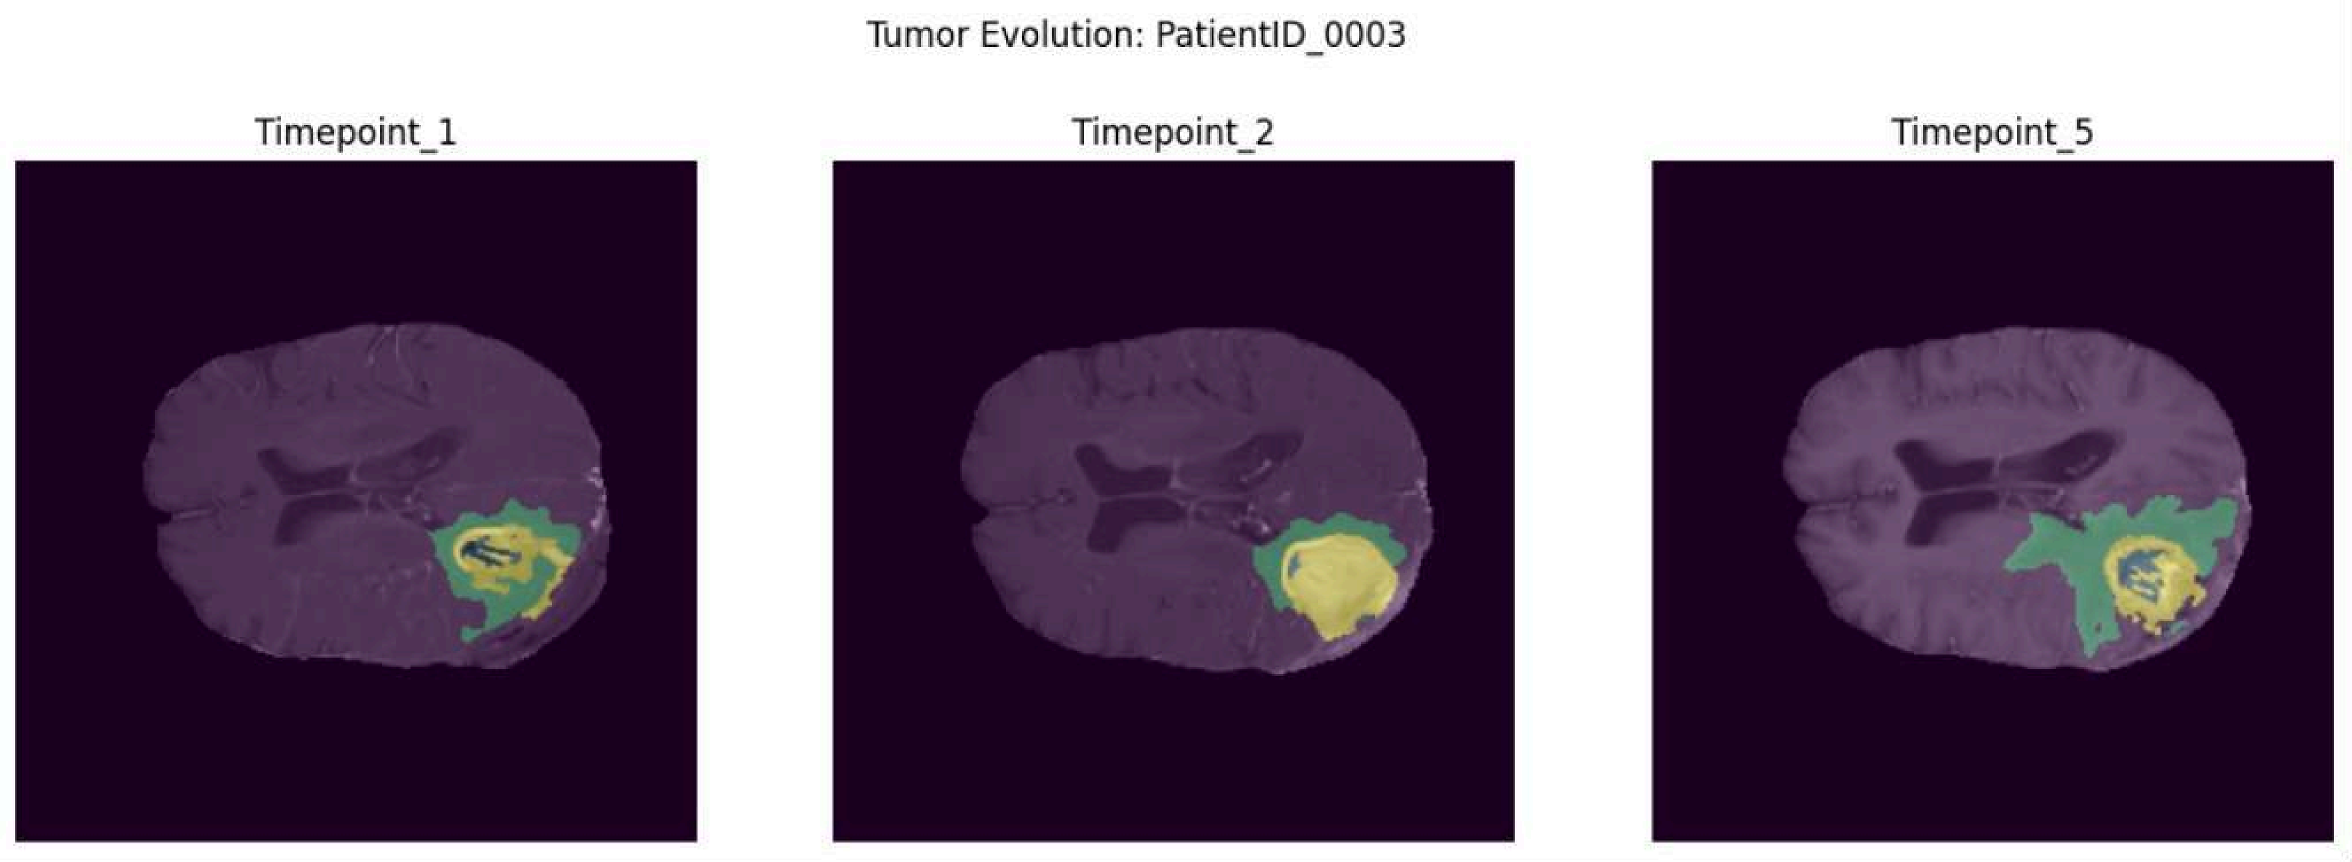

Sources & Ethics

Dataset Source:
Yaseen, D., Garrett, F., Gass, J., Greaser, J., Isufi, E., Layfield, L. J., Nada, A., Porgorzelski, K., Sinclair, J., Tahon, N. H. M., & Thacker, J. (2025). University of Missouri Post-operative Glioma Dataset (MU-Glioma-Post) (Version 1) [Data set]. The Cancer Imaging Archive. https://doi.org/10.7937/7K9K-3C83

The dataset used in this study was obtained from an authorized source, and permission was granted for its use in research and academic purposes. All patient data were handled in accordance with ethical guidelines, ensuring privacy, confidentiality, and responsible use of medical information.

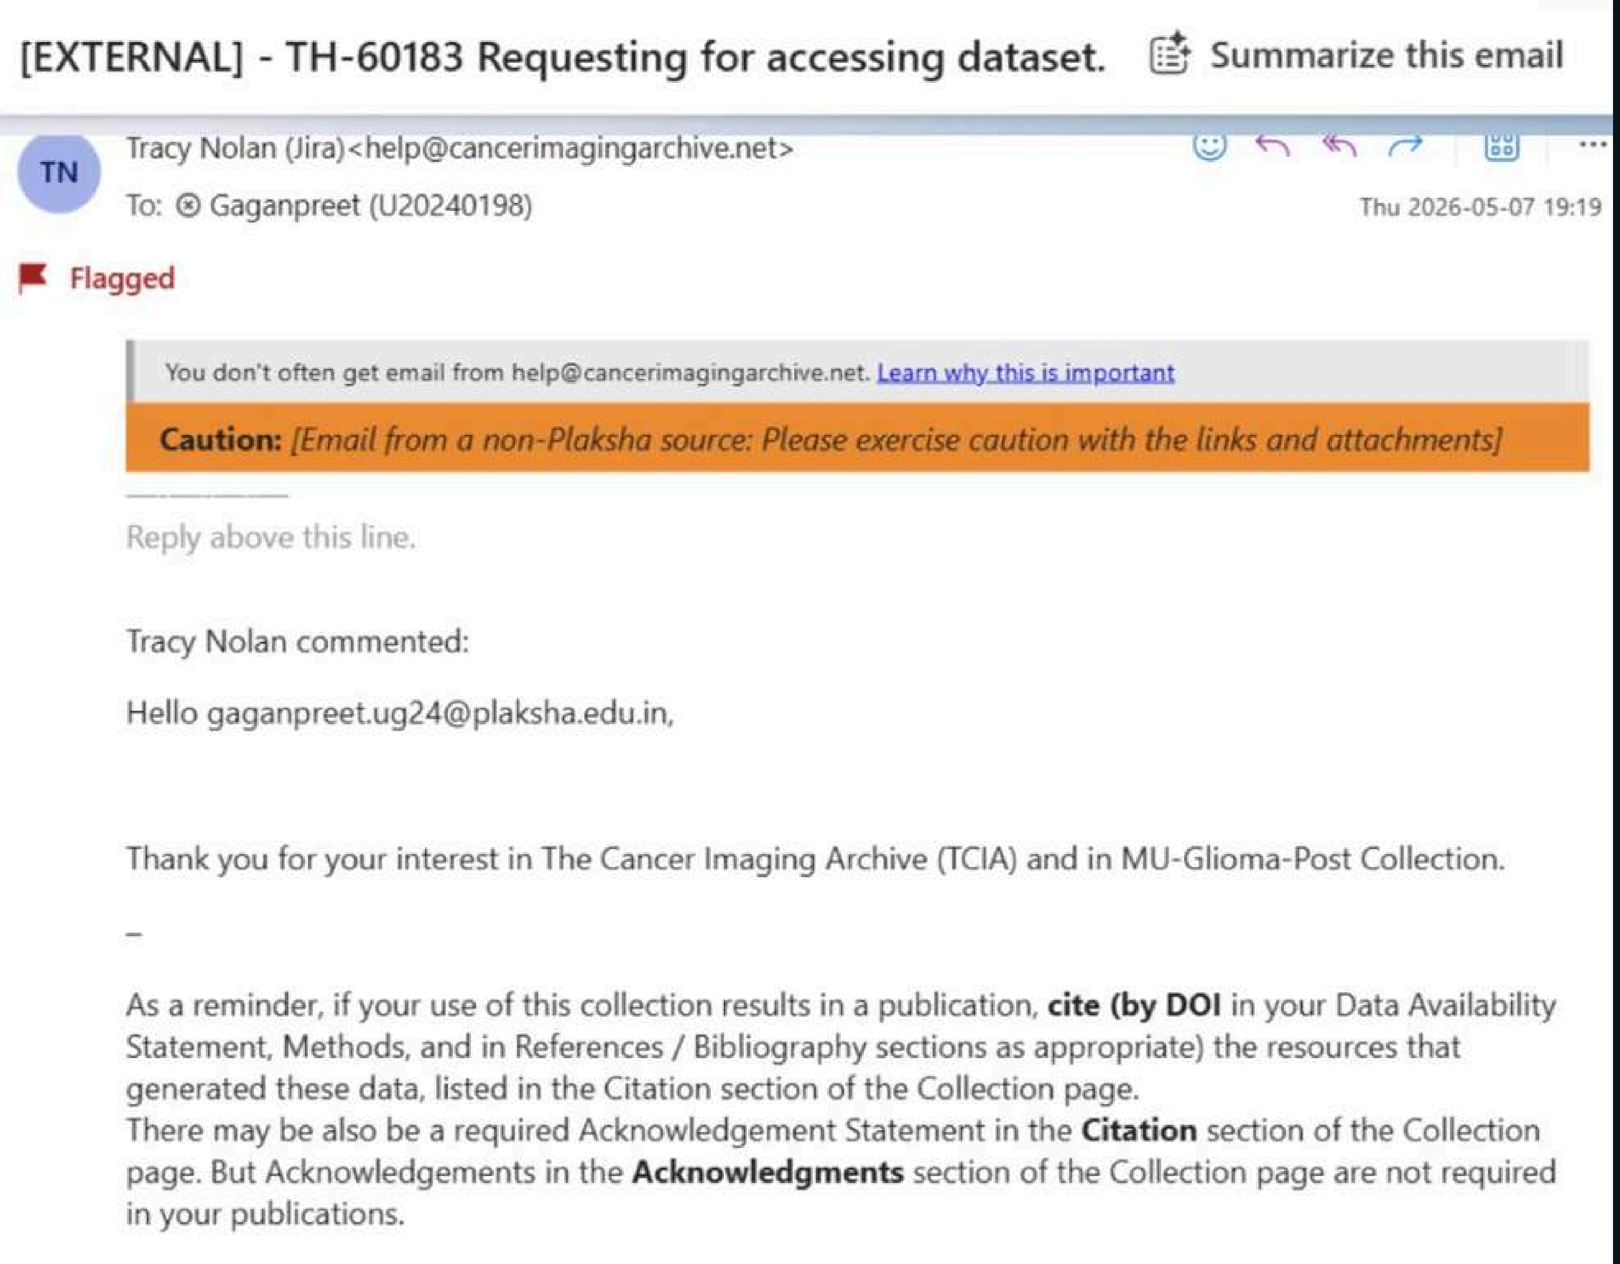

[EXTERNAL] - TH-60183 Requesting for accessing dataset.

Tracy Nolan (Jira)<help@cancerimagingarchive.net>
To: Gaganpreet (U20240198)
Thu 2026-05-07 19:19

You don't often get email from help@cancerimagingarchive.net. Learn why this is important

Caution: *[Email from a non-Plaksha source: Please exercise caution with the links and attachments]*

Reply above this line.

Tracy Nolan commented:

Hello gaganpreet.ug24@plaksha.edu.in,

Thank you for your interest in The Cancer Imaging Archive (TCIA) and in MU-Glioma-Post Collection.

–

As a reminder, if your use of this collection results in a publication, **cite (by DOI** in your Data Availability Statement, Methods, and in References / Bibliography sections as appropriate) the resources that generated these data, listed in the Citation section of the Collection page.
There may be also be a required Acknowledgement Statement in the **Citation** section of the Collection page. But Acknowledgements in the **Acknowledgments** section of the Collection page are not required in your publications.

Feature Preprocessing (Clinical)

1. Remove Unnecessary Features (e.g. Race)
2. Handle Missing Values
3. Separate Feature Types (Numeric & Categorical)
4. Encode Categorical Variables
5. Feature and Target Selection
6. Feature Scaling

Used PCA for dimensionality Reduction

Feature Preprocessing (MRI)

1. Build Patient Timepoint Index
2. MRI Intensity Normalization
3. Convert Mask tumors into 4 chanels
4. Build the 8-Channel Tensor (4 MRI & 4 tumor channels)
5. Tumor-Centered 3D Patch Extraction
6. Patch Padding
7. Scan Padding

plots of feature engineering :D

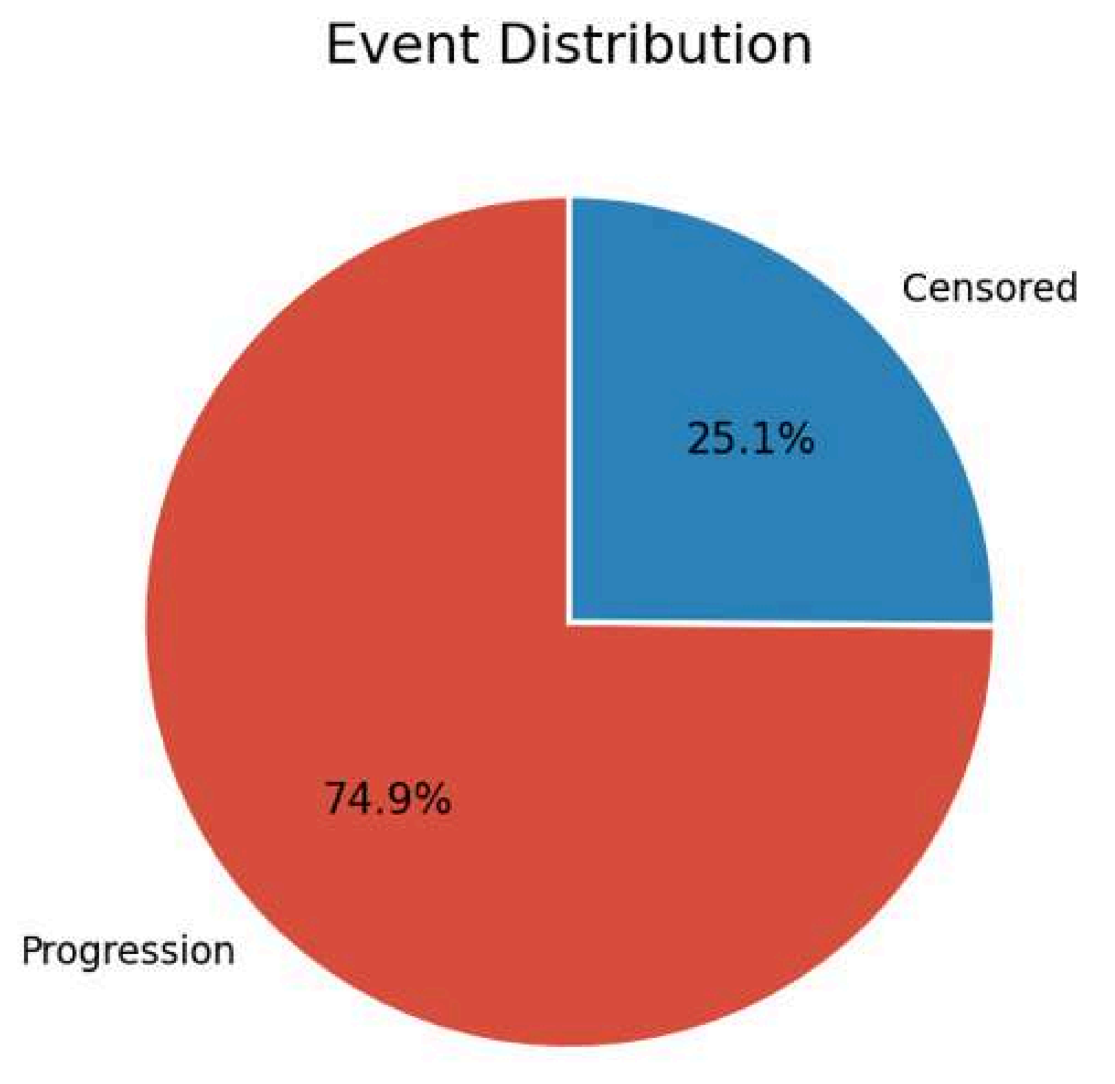

Majority of the patients belong to progression group, and smaller portion of cases are censored

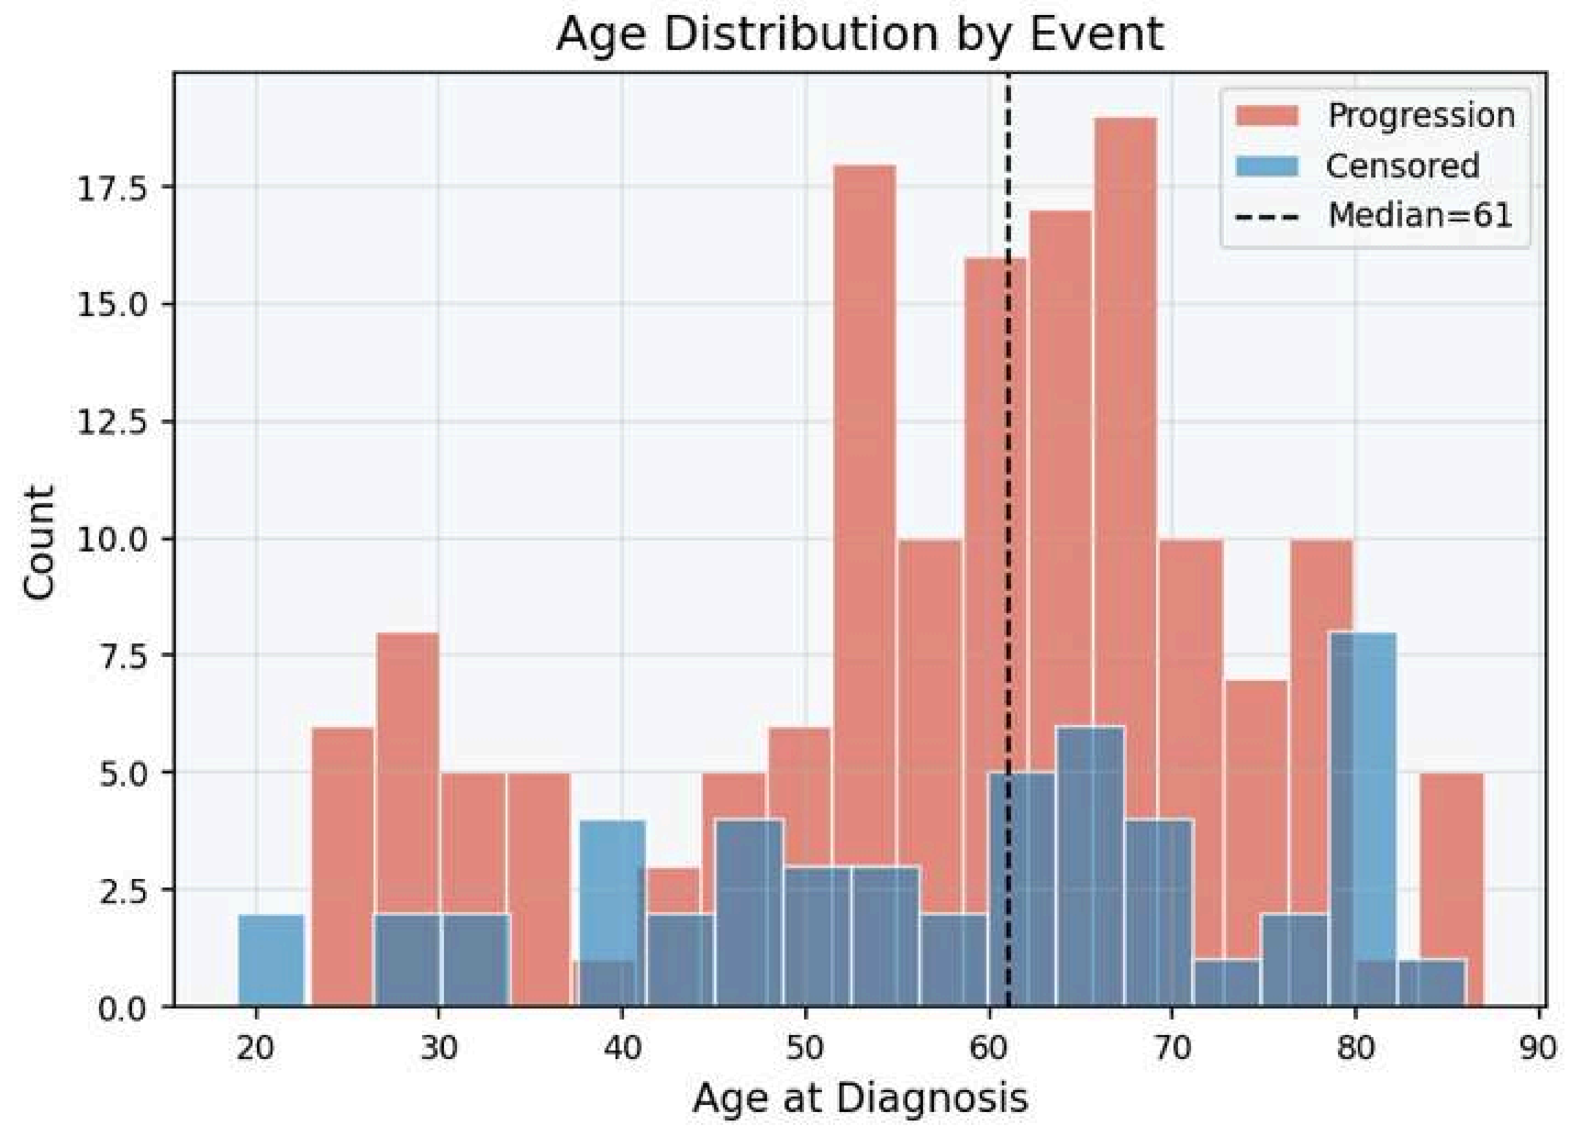

Most patients are diagnosed between 50-70 years.
Median age: 61 years.
Porgression events are frequent in older patients.

plots of feature engineering :D

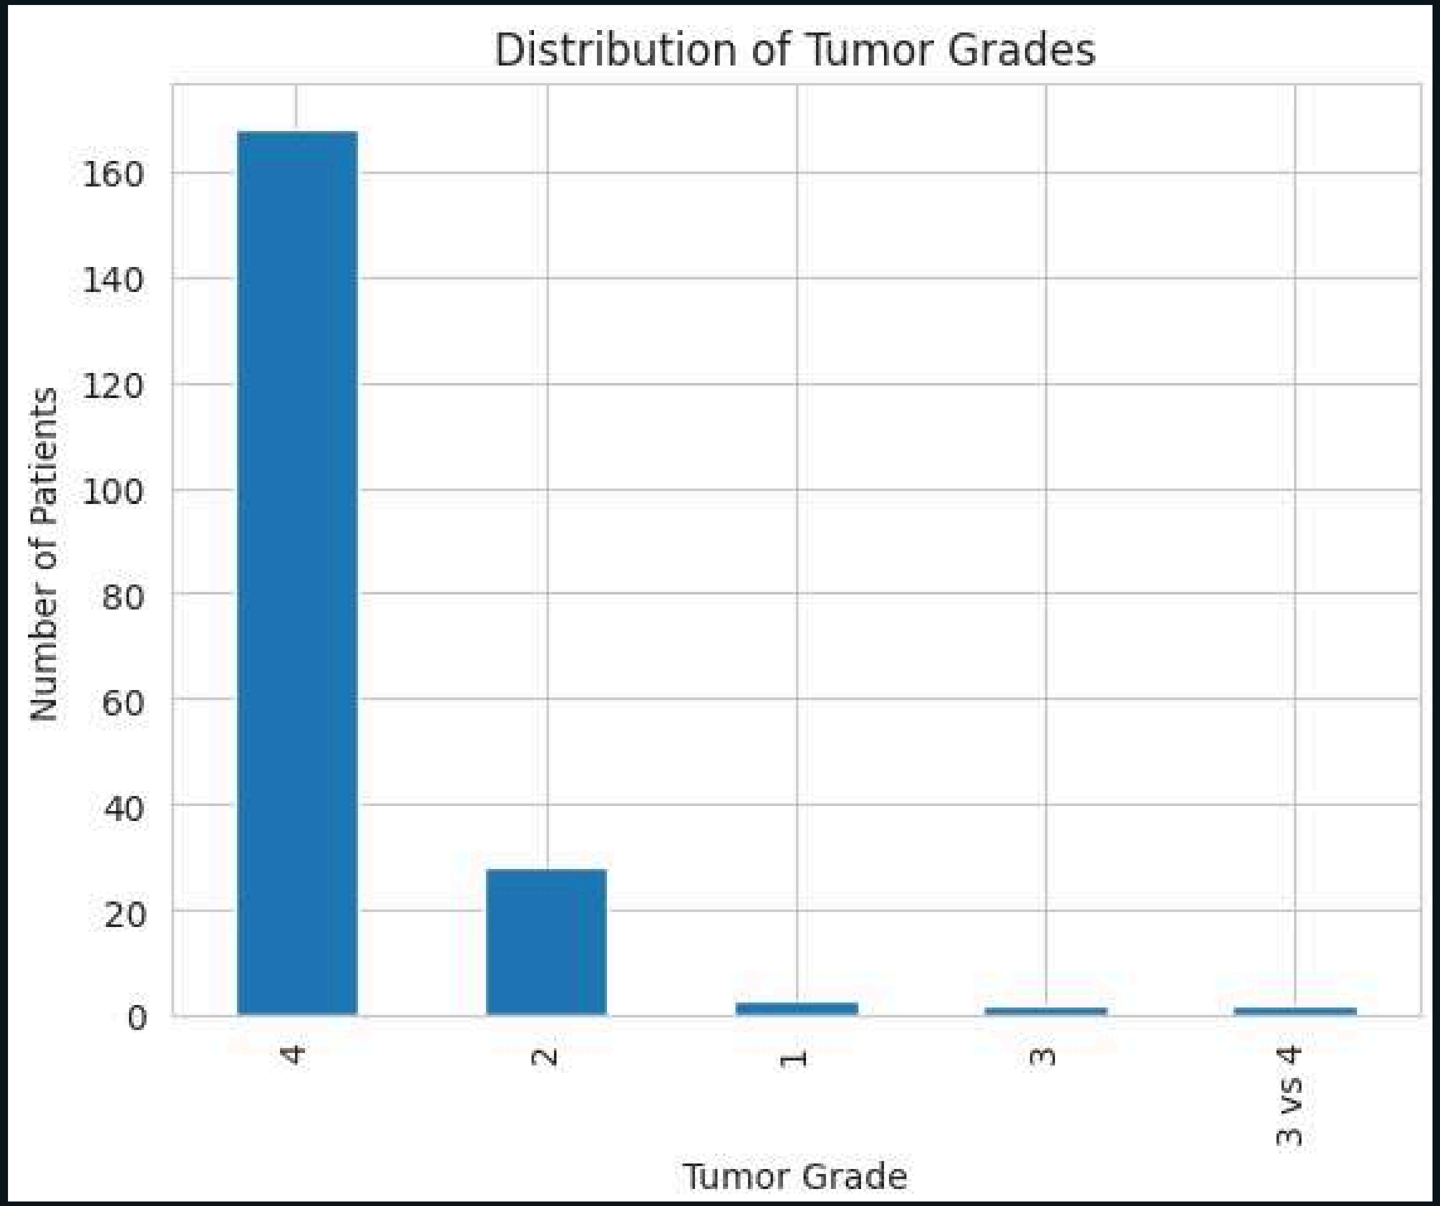

Classification of tumors by the Tumor Grade. Majority of the patients belong to Grade 4 tumors.

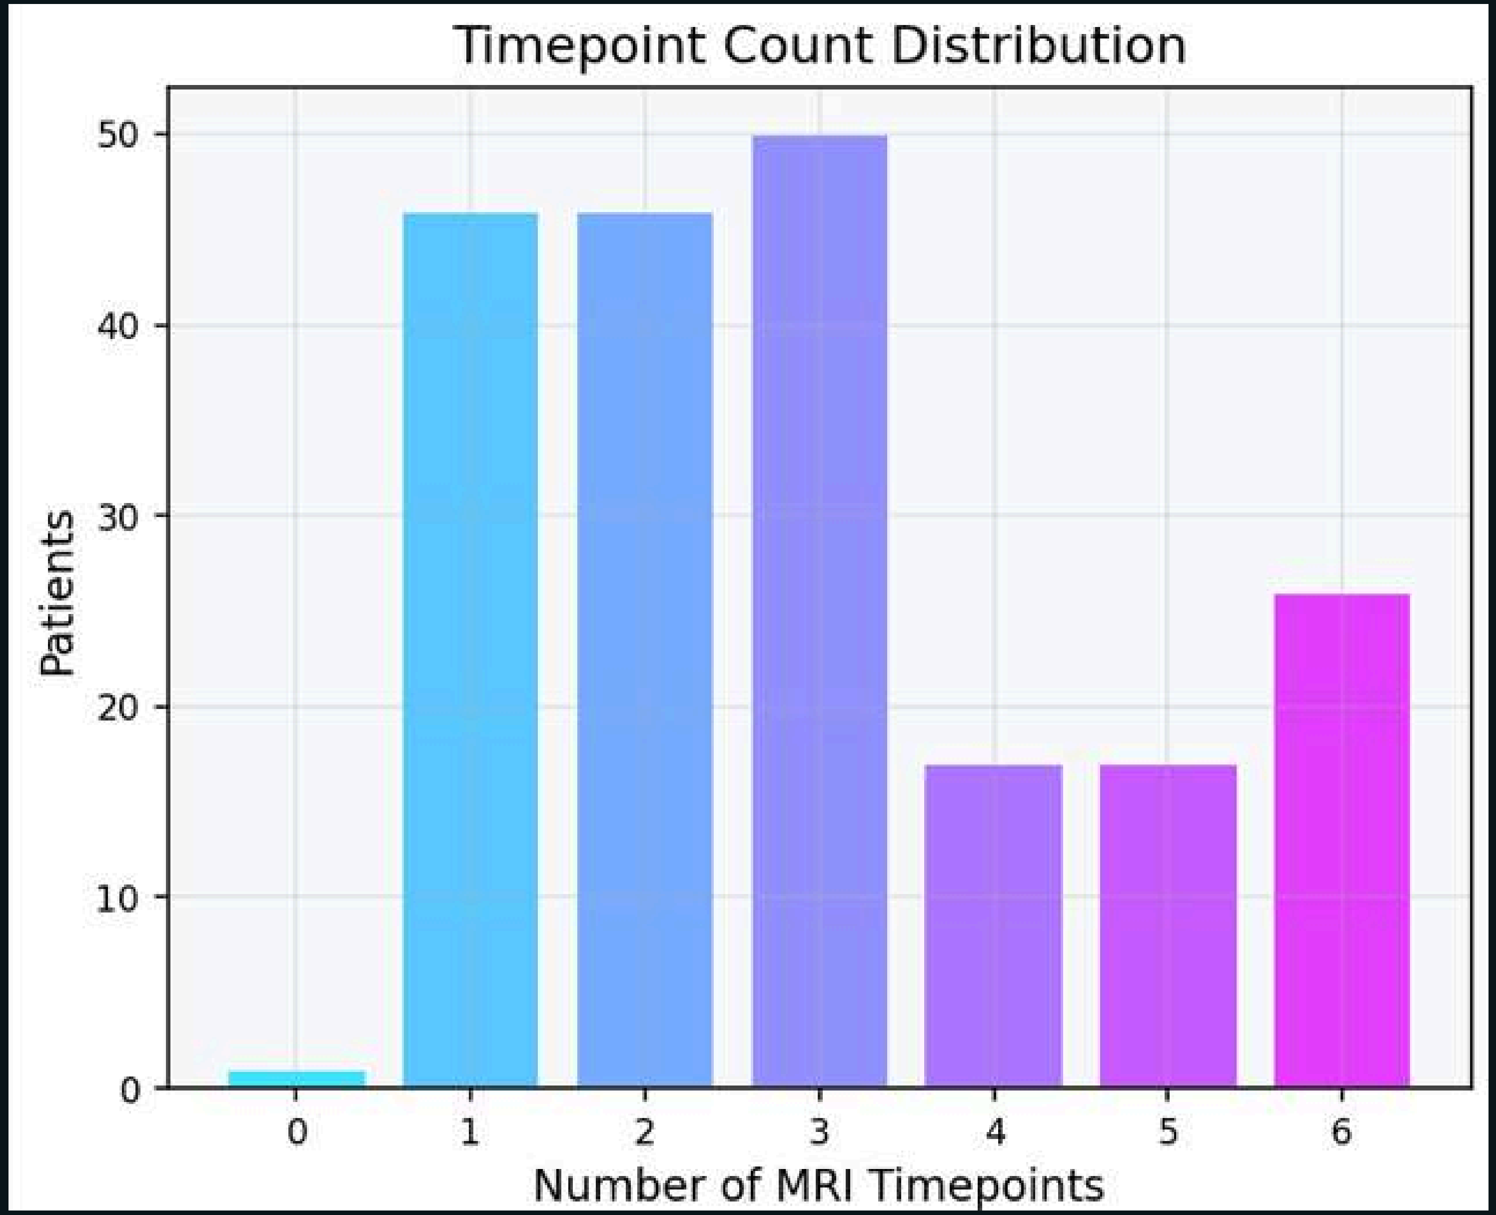

Most pateients have 1-3 MRI timepoints available. Fewer patients contain long term follow up scans, indicating variability of MRI availability accross patients.

plots of feature engineering :D

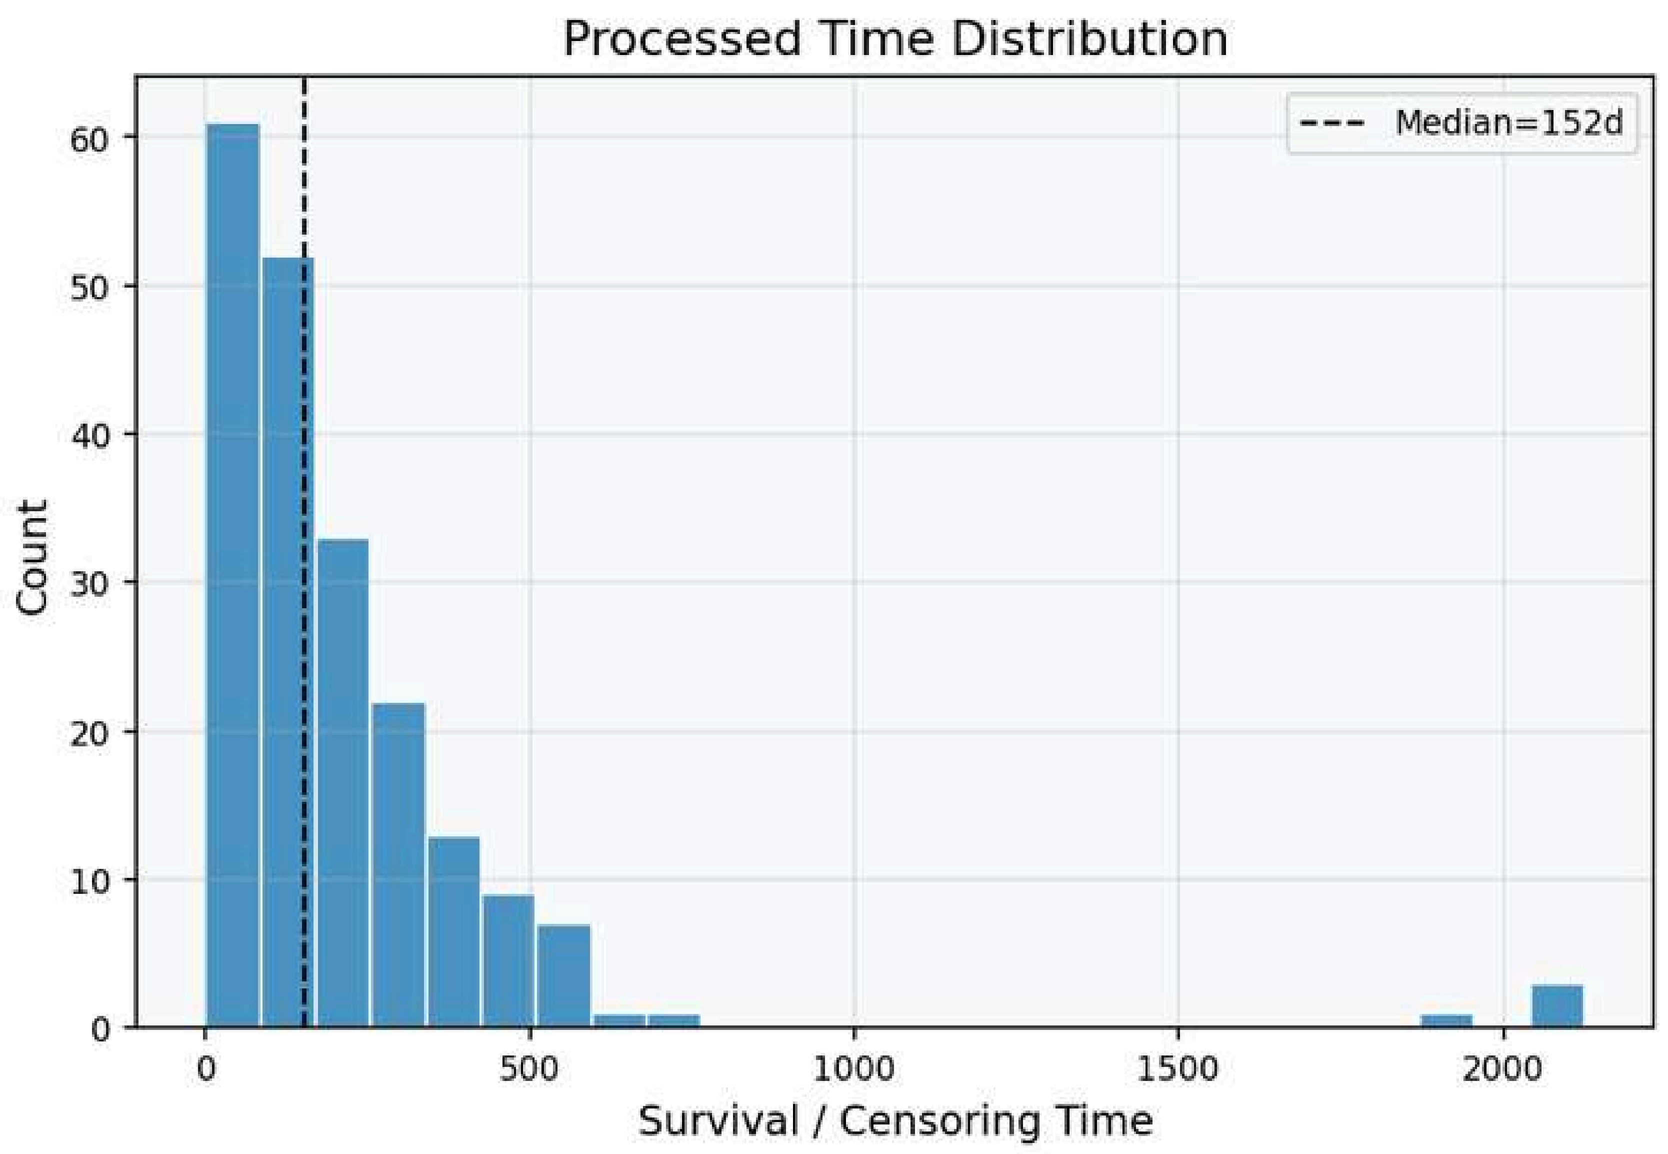

Most survival times are concentrated below 500 days, with a median of 152 days.

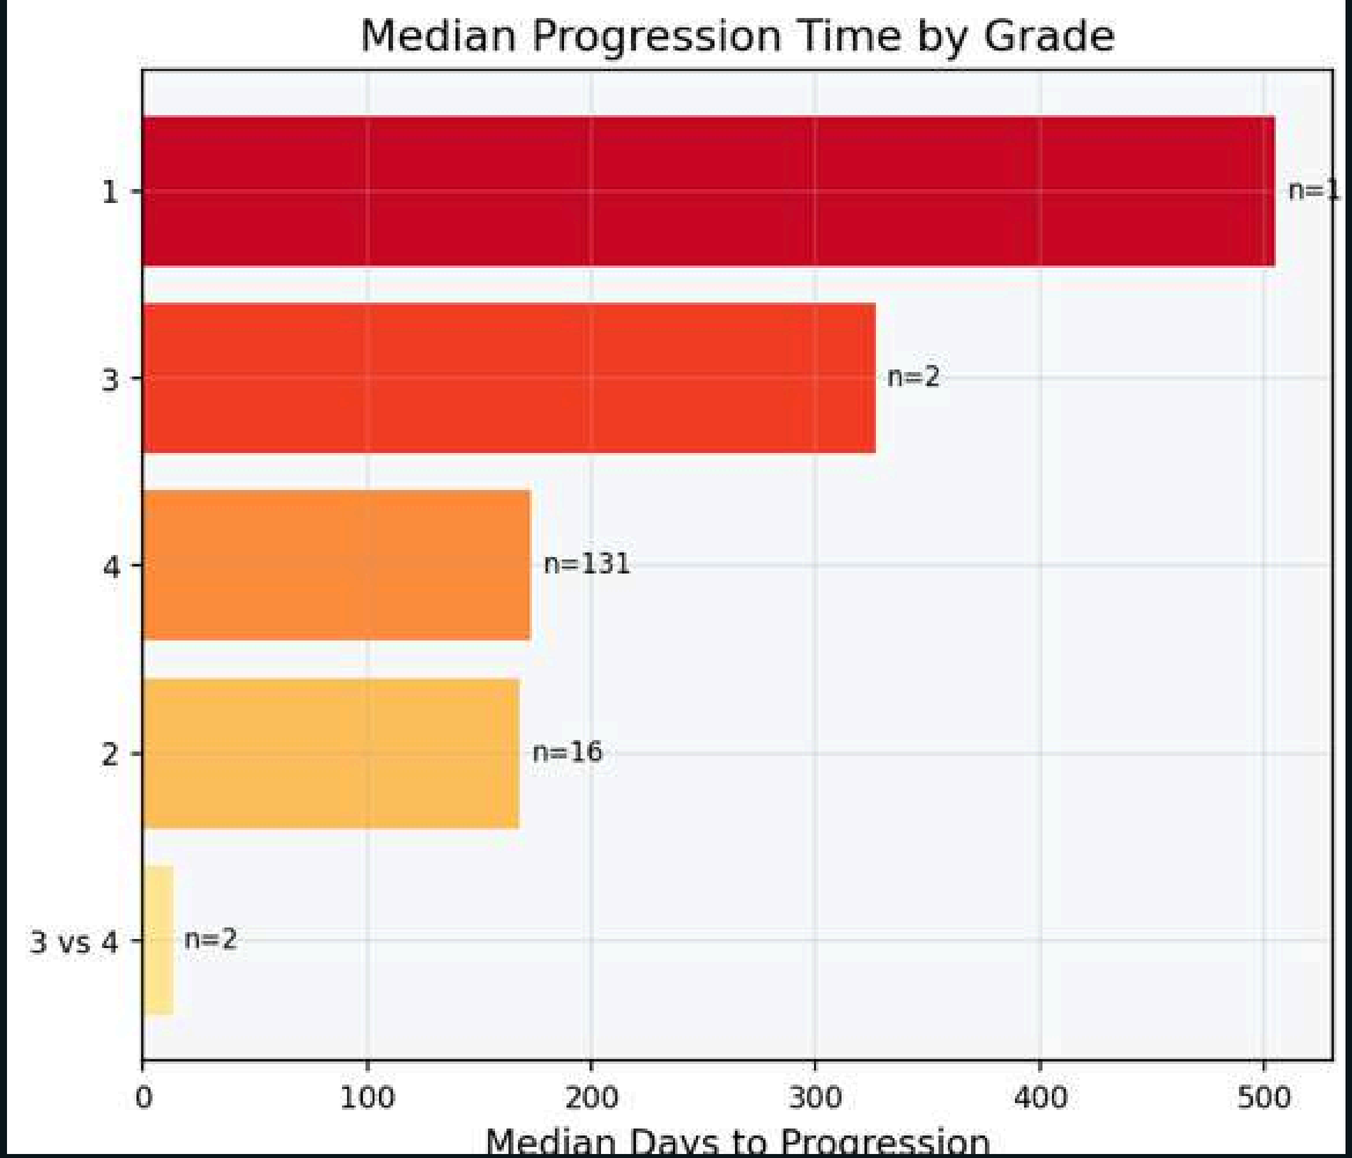

Lower tumor grades show longer median progression times, while Grade 4 tumors have comparatively shorter progression durations.

METHODOLOGY

Custom CNN (3.6 million parameters)

Block 1

Conv3D: 8 → 32 channels
BatchNorm + ReLU + MaxPooling
Extracts basic features like edges and textures

Block 2

Conv3D: 32 → 64 channels
BatchNorm + ReLU + MaxPooling
Learns local tumor patterns and structures

Block 3

Conv3D: 64 → 128 channels
BatchNorm + ReLU + MaxPooling
Captures complex tumor morphology and growth patterns

Block 4

Conv3D: 128 → 256 channels
BatchNorm + ReLU + Adaptive Average Pooling
Generates compact global MRI features

Fully Connected Layer

Feature Vector: 256 → 512
Produces a 512-dimensional latent embedding

What we found from other papers:

1

ResNet-18

- Extracts tumor shape, texture, edema, and volumetric patterns from multi-modal MRI scans
- Learns deep spatial representations for each MRI timepoint
- Residual connections improve feature learning and reduce overfitting
- Generates 512-dimensional embeddings used for downstream temporal modeling.

2

U-Net

- Segments tumor regions: necrosis, edema, enhancing tumor core
- Improves focus on clinically important regions before feature extraction
- Uses encoder–decoder architecture with skip connections for precise localization
- Dice coefficient and Mean IoU commonly used for evaluation

3

MobileNet-Small

- Lightweight CNN with fewer parameters and faster inference speed
- Performs well on limited medical datasets with reduced computational cost
- Achieved strong classification accuracy among compared architectures
- Suitable for deployment in resource-constrained healthcare environments

---

Brain tumor detection	Ensemble Technique for Brain Tumor Analysis	Employing deep learning and transfer learning for accurate brain tumor detection
Parameters: 11.7 Million	Parameters: 31 Million	Parameters: 2.5 Million
Cross Entropy Loss	Categorical Cross entropy Loss	Cross Entropy Loss
C-Index (0.45-0.52) Training loss was high	Bad C-index (0.35-0.45)	Implemented in Proposed Model

MobileNet Architecture

1. Light weight CNN used for MRI feature extraction
2. Modified input layer: 3 → 8 channels
3. Used depthwise separable convolutions for efficient computation
4. Pointwise convolutions combine multimodal MRI features
5. Includes Inverted Residual Blocks
6. Uses squeeze and excitation to focus on important tumor regions
7. Hard-swish activation for optimization and feature learning
8. Extracts compact spatial embeddings from MRI slices
9. Final vector passed to LSTM

LSTM & Prediction Head Flow

A deep learning pipeline that processes sequential MRI scans to predict tumour progression risk over time.

1

A) Sequential MRI Embeddings

- Receives 512-D embeddings from MobileNet encoder
- One embedding generated for each MRI timepoint
- Forms longitudinal sequence: (B, T, 512)

2

B) Input Dropout (0.2)

- Applied before LSTM processing
- Reduces overfitting on limited patient data
- Improves generalization performance

3

C) Long Short-Term Memory (LSTM)

- Learns temporal tumor evolution across scans
- Captures progression patterns such as:
  - Tumor growth
  - Stabilization after treatment
  - New enhancing regions
  - Progression recurrence
- Stores long-term temporal information

4

D) Hidden State Representation

- Final hidden state summarizes full patient history
- Produces compact temporal embedding: (B, 256)
- Represents overall progression trajectory

5

E) Prediction Head

- Fully connected layers convert temporal embedding → risk score
- Architecture: 256 → 128 → ReLU → Dropout → 1
- Generates patient-level progression risk

6

F) Cox Survival Prediction

- Output = Log-risk score (not probability)
- Higher score → Higher predicted tumor progression risk
- Cox Loss ranks patients based on survival risk over time

Summary

Input: (B, T, 512) | **LSTM Output:** (B, 256) | **Final Output:** (B, 1)

Performance Metrics Used:

Metric	What it says
C-Index	measures how well our model correctly ranks brain tumor patients based on their predicted survival risk and actual survival time
IBS	The IBS (Integrated Brier Score) measures how accurate our model's predicted survival probabilities are over time
TD-AUC	measures how well our model distinguishes high-risk and low-risk patients at different time points during survival prediction

Performance Metrics:

Metric	Value Obtained	Values in Papers
C-Index	0.6894	0.65-0.80
IBS	0.28	0.12 - 0.22
TD-AUC	0.92	0.7-0.93

Performance Metrics:

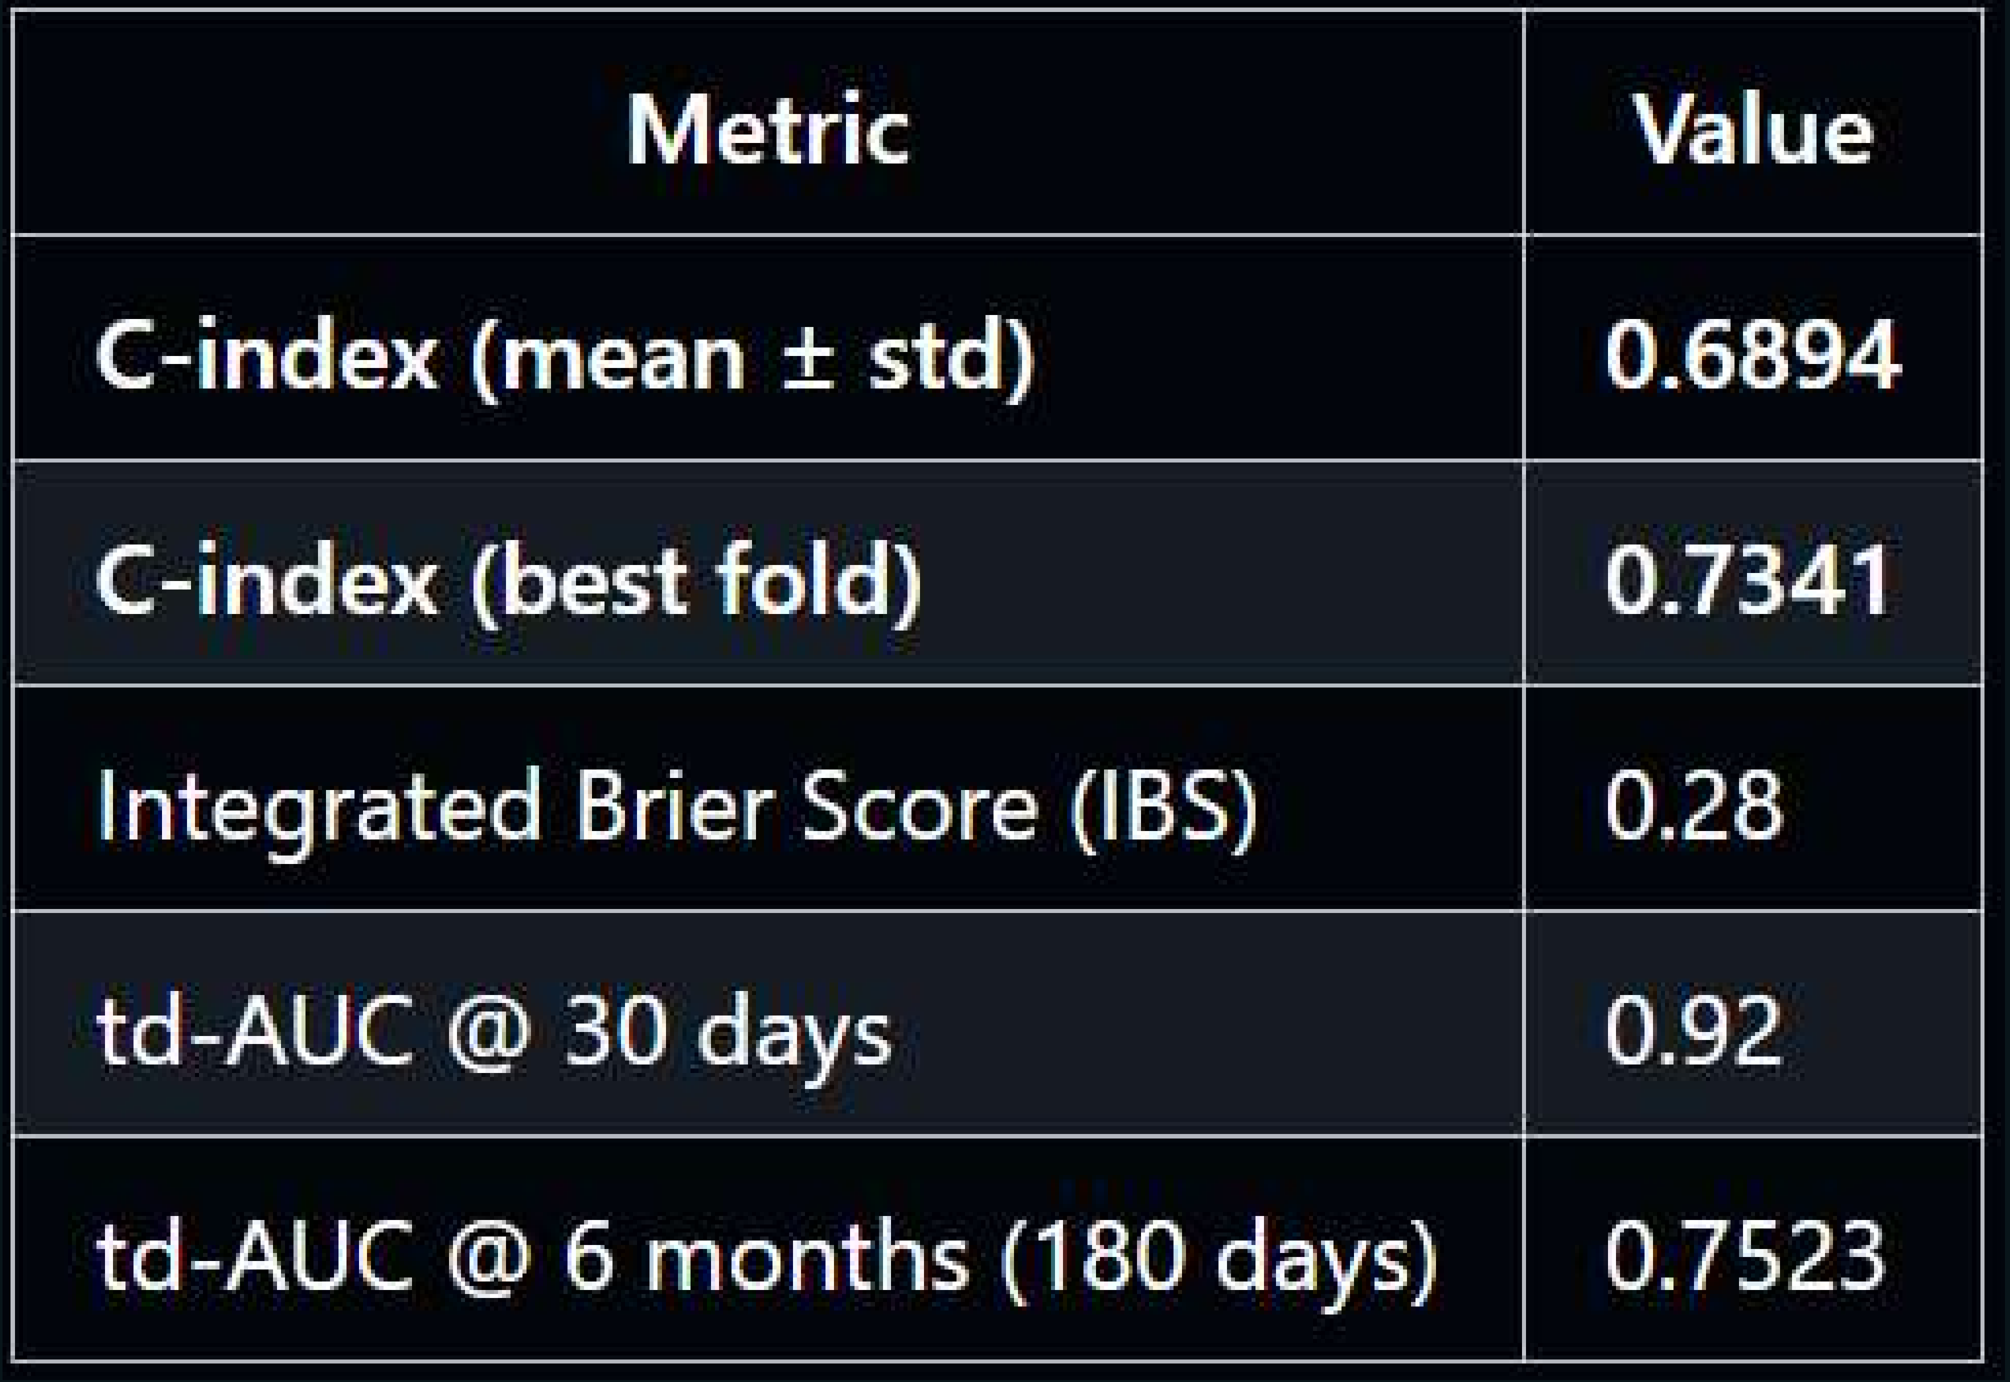

Metric	Value
C-index (mean ± std)	**0.6894**
C-index (best fold)	**0.7341**
Integrated Brier Score (IBS)	0.28
td-AUC @ 30 days	0.92
td-AUC @ 6 months (180 days)	0.7523

Implementation Challenges & Solutions

Challenge	Our Solution
Limited Dataset — 203 patients	Gaussian noise, intensity scaling + Dropout + L2 regularization + 3-fold cross validation
Variable Scan Sequences — 1 to 6	Zero-pad to 6 timepoints + attention mask (1=real, 0=padded time points) + day-from-diagnosis timestamps as auxiliary input

THANKYOU

Proposed Machine Learning Methodology

The proposed machine learning pipeline combines spatial and temporal modeling to analyze longitudinal MRI data. Multi-modal MRI scans and tumor segmentation masks are preprocessed and stacked as multi-channel inputs. A 3D ResNet-18 network extracts spatial tumor features such as shape, volume, and spatial heterogeneity from each MRI scan. These spatial embeddings are then passed into an LSTM network that models how the tumor evolves across multiple timepoints. Finally, a prediction head estimates the probability that tumor progression will occur within a given time horizon.

Three Key Components

1

Spatial Feature Extraction (3D ResNet-18)

A 3D convolutional neural network (3D ResNet-18) will be used to extract spatial features from each MRI scan. The CNN processes the multi-modal MRI volume and tumor masks to learn representations that capture tumor shape, volume, spatial heterogeneity, and structural patterns.

2

Temporal Modeling (LSTM)

Since each patient has multiple MRI scans over time, a Long Short-Term Memory (LSTM) network will be used to model the temporal evolution of tumor features. The LSTM processes the sequence of feature embeddings extracted from the CNN and learns patterns such as tumor growth, stabilization, or infiltration across follow-up scans.

3

Prediction Head

The output of the LSTM represents a patient-level feature representation. A fully connected prediction head will convert this representation into a tumor progression risk estimate. The model will estimate the probability that tumor progression will occur within a specified future time horizon.

Why This Approach Works

What We Need	Why It Matters	Our Algorithm
Spatial — tumor shape, volume, heterogeneity	Tumor structure at each timepoint tells current disease state	3D ResNet-18 CNN
Temporal — how tumor evolves across scans	Progression is a process, not a snapshot	LSTM
Clinical — demographics, molecular markers, treatment	Imaging alone misses biological context	Feature fusion in prediction head