


# BTReport: A Framework for Brain Tumor Radiology Report Generation with Clinically Relevant Features

Juampablo E. Heras Rivera<sup>\*1</sup> 


JEHR@UW.EDU

Dickson T. Chen<sup>\*1</sup> 

DTCHEN19@UW.EDU

Tianyi Ren<sup>1</sup> 


TR1@UW.EDU

Daniel K. Low<sup>2</sup> 


DALOW@UW.EDU

Jacob Ruzevick<sup>2</sup>

RUZEVICK@UW.EDU

Asma Ben Abacha<sup>3</sup> 

ABENABACHA@MICROSOFT.COM

Alberto Santamaria-Pang<sup>3,4</sup> 

ALBERTO.SANTAMARIAPANG@MICROSOFT.COM

Mehmet Kurt<sup>1</sup> 

MKURT@UW.EDU

<sup>1</sup> *University of Washington*

<sup>2</sup> *University of Washington School of Medicine*

<sup>3</sup> *Microsoft Health AI*

<sup>4</sup> *The Johns Hopkins University School of Medicine*

**Editors:** Accepted for publication at MIDL 2026

## Abstract

Recent advances in radiology report generation (RRG) have been driven by large paired image-text datasets; however, progress in neuro-oncology RRG has been limited due to a scarcity in open paired image-report datasets. Here, we introduce BTReport, an open-source framework for brain tumor RRG that constructs natural language radiology reports using reliably extracted quantitative imaging features. Unlike existing approaches that rely on general-purpose or fine-tuned vision-language models for both image interpretation and report composition, BTReport performs deterministic feature extraction of clinically-relevant features, then uses large language models only for syntactic structuring and narrative synthesis. By separating RRG into deterministic feature extraction and report generation stages, synthetically generated reports are completely interpretable and contain reliable numerical measurements, a key component lacking in existing RRG frameworks. We validate the clinical relevance of BTReport-derived features, and demonstrate that BTReport-generated reports more closely resemble reference clinical reports when compared to existing baseline RRG methods. To further research in neuro-oncology RRG, we introduce BTReport-BraTS, a companion dataset that augments BraTS imaging with synthetic radiology reports generated with BTReport, and BTReview, a web-based platform for validating the clinical quality of synthetically generated radiology reports. Code for this project can be found at: <https://github.com/KurtLabUW/BTReport>.

**Keywords:** Brain MRI, Radiology report generation, VASARI, Midline shift, Open dataset, Multimodal learning, Neuro-oncology

## 1. Introduction

Radiology is a medical specialty that employs a variety of imaging modalities (e.g., X-ray, computed tomography (CT), multi-parametric magnetic resonance imaging (mpMRI) for

---

\* Contributed equally

the detection and monitoring of human disease. The radiology report contains a detailed summary of imaging findings, providing insights into a patient’s condition crucial for diagnosis and clinical decision making. With a growing aging population, the demand for radiology services is expected to increase between 16.9% to 26.9% by 2055, while attrition in radiology also continues to increase (Christensen et al., 2025). As the gap between physician workload and available workforce widens, assisted radiology report generation (RRG) is positioned to help address this unmet clinical need.

RRG leverages advances in artificial intelligence (AI) to extract quantitative imaging markers from raw unstructured data in an automated manner. In clinical workflows, RRG promises to improve data quality, repeatability, factual completeness, and timeliness of radiology reporting. The adoption of vision-language models (VLMs) in RRG has led to substantial advances, allowing models to jointly reason over medical images and textual findings (Liu et al., 2019; Chen et al., 2020; Sellergren et al., 2025). These developments have been primarily driven by large-scale chest X-ray datasets such as MIMIC-CXR (Johnson et al., 2019) and IU X-ray (Demner-Fushman et al., 2016), which provide over 300,000 chest X-ray images paired with clinical reports. However, accessible image-report datasets across other radiology specialties are limited. For instance, in neuro-oncology, large neuroimaging datasets are made openly available through efforts such as BraTS (de Verdier et al., 2024), but paired text reports for VLM training remain lacking.

Advances in computer vision have automated glioblastoma (GBM) segmentation from mpMRI, enabling reliable quantification of the contrast-enhancing lesion, necrotic core, and peritumoral edema (Menze et al., 2014). However, segmentations alone lack the clinical context needed to evaluate the impact of tumor on the surrounding brain environment. Neuro-radiologists supplement their interpretation with additional imaging-derived evidence, such as subregion involvement and midline shift (MLS) measurements, among others. Together, these descriptors provide comprehensive brain tumor characterization, and improve risk stratification, treatment planning, and outcome prediction. Neuro-oncology RRG models should reliably include these features to more closely align with radiologist expectations.

Here we introduce BTRReport, a two-stage framework for neuro-oncology RRG grounded in clinically relevant quantitative features. BTRReport first deterministically extracts imaging markers from mpMRI including patient metadata, VASARI (Visually AcceSAbLe Rembrandt Images) features, and automated 3D midline shift measurement. These markers are then provided as structured inputs to large language models (LLMs) for clinical reasoning and synthetic report generation. Because quantification is performed algorithmically upstream, BTRReport enables measurement-grounded reporting without requiring task-specific fine-tuning of vision encoders or VLMs. This design addresses a key challenge in neuro-oncology RRG: clinical reports frequently include quantitative measurements such as lesion size, volume, and midline shift, which may not be reliably extracted by generic vision encoders without explicit training for measurement and precise spatial reasoning (Chen et al., 2024). By grounding report generation in deterministically extracted imaging features, our framework is directly interpretable and reduces the likelihood of critical detail omission (Wu et al., 2025).

Our contributions are as follows: **a)** a scalable brain tumor mpMRI radiology report generation framework driven by deterministically extracted neuroimaging features (Section 4.1), **b)** a robust, interpretable 3D midline shift (MLS) estimation algorithm (Section 4.3),

c) clinical validation of BTReport-derived features demonstrated with semantic clustering of radiology concepts from reference reports, and retrospective modeling of overall survival (Section 4.5.2), and d) release of BTReport-BraTS, an open-source image-report companion dataset augmenting BraTS cases with clinically grounded anatomical and pathological descriptors (Appendix G).

## 2. Related Work

A variety of approaches have been proposed for the task of image-paired RRG and generally fall into one- and two-stage frameworks. The leading paradigm for RRG involves training monolithic VLM foundation models to extract image features and generate reports in one step, such as in MedGemma (Sellergren et al., 2025) and MedPaLM-2 (Singhal et al., 2023). Approaches in neuro-oncology following this paradigm include TextBraTS (Shi et al., 2025), which directly prompts GPT-4 models (Achiam et al., 2023) with videos of 2D mpMRI axial slices alongside tumor segmentation masks derived from FLAIR imaging. Radiologists refine these annotations into structured textual labels, which are then used to guide tumor segmentation models toward clinically relevant regions, improving segmentation performance. Although these annotations improved segmentation performance, they do not contain quantitative measurements typically reported by radiologists.

Other image-paired RRG frameworks adopt a two-stage design. For example, *From Segmentation to Explanation (Seg-to-Exp)* (Valerio et al., 2025) first establishes descriptive tumor-ROI relationships based on the co-registration of tumor segmentation masks to an anatomical atlas. These relationships are extracted as structured features that describe the percent overlap of the brain tumor with ROIs it occupies, and then provided to an LLM for report generation. The generated reports provide details about regional tumor impact and offer insightful clinical implications based on the functional role associated with identified tumor-ROI relationships. However, generated reports do not resemble the standard narrative structuring of radiology reports authored by fellowship trained neuro-radiologists. Additionally, *Seg-to-Exp* reports do not describe additional clinically-relevant imaging observations beyond anatomical tumor location.

*AutoRG-Brain* (Lei et al., 2024) also adopts a two-stage framework. In the first stage, anomaly ROI masks are generated and mapped to specific ROIs based on anatomical segmentations. In the second stage, a fine-tuned VLM takes embedded anomaly ROI masks as visual prompts for RRG. This approach uses deterministic features to focus the VLM, generating region-specific findings grounded in areas of the image where anomalies were detected. Based on this design, generated reports are largely descriptive, with findings expressed through signed magnitude operators (e.g., “enlarged”, “high signal”) that encode comparative attribute relationships between detected anomalies and their anatomical location. Importantly, describing an imaging abnormality in relation to anatomical context does not, by itself, establish the clinical context required for higher-order clinical interpretation. As a result, observations from generated reports are not framed in a clinically meaningful way, limiting their utility in informing downstream clinical decision making. Furthermore, reference reports used for fine-tuning contain a limited set of clinically-relevant features.

The framework proposed in *RadGPT* (Bassi et al., 2025) first deterministically extracts clinically-relevant features from CT scans of abdominal tumors, then uses LLMs for syntactic

structuring. This approach produces radiology reports that more closely resemble reference reports when compared to end-to-end report generation models. However, it remains unclear whether a framework similar to *RadGPT* can be applied for neuro-oncology RRG based on deterministically extracted features from mpMRI.

### 3. Data

This study uses two complementary datasets to support radiology report generation and validation of BTRReport-derived features. First, the HuskyBrain dataset is a retrospective cohort of fully de-identified pre-operative mpMRI studies paired with radiologist-authored reports, serving as the primary resource for training and evaluation of neuro-oncology RRG frameworks. The second dataset is a subset of BraTS’23 cases, which enable survival analyses with imaging and clinical metadata. Together, these datasets support the evaluation of synthetic report clinical quality and the use of deterministic neuroimaging features for survival outcome prediction.

#### 3.1. HuskyBrain Dataset

We collected pre-operative mpMRI scans and radiology reports from a retrospective cohort of GBM patients (n=184) treated at the University of Washington Medical Center (UWMC), an academic medical institution serving patients across the WWAMI region (Washington, Wyoming, Alaska, Montana, and Idaho). Inclusion criteria were: 1) confirmed histopathologic diagnosis of GBM; 2) availability of pre-operative mpMRI sequences including T1, T2, T1c, and T2-FLAIR; and 3) a corresponding pre-operative diagnostic radiology report. Images were pre-processed using CaPTk (Davatzikos et al., 2018; Pati et al., 2020), following BraTS 2017-2023 procedures: DICOM to NIfTI conversion, SRI24 co-registration, 1 mm isotropic resampling, and skull-stripping. Reference reports were authored by fellowship-trained radiologists and selected as the clinical reference standard. For each case, the HuskyBrain dataset contains de-identified MRI sequences, radiologist-authored reports stripped of protected health information, and tumor masks.

#### 3.2. Survival Analysis Dataset

From the BraTS’23 dataset (Adewole et al., 2023), a smaller cohort of mpMRI cases (n=461) was used for survival analyses. Cases were selected based on the availability of five minimum metadata entries: (1) age at initial diagnosis, (2) biological sex, (3) confirmed methylation status of O6-methylguanine-DNA methyltransferase (*MGMT*), (4) mutation status of isocitrate dehydrogenase 1 (*IDH1*), and (5) overall survival or equivalent survival quantification representing the number of days between radiological diagnosis and reported days to known death. We collected these demographic and genomic features from multiple publicly available collections of GBM cases, including those from the University of California San Francisco (UCSF-PDGM) (Calabrese et al., 2022), University of Pennsylvania (UPenn-GBM) (Bakas et al., 2021), Clinical Proteomic Tumor Analysis Consortium (CPTAC-GBM) (National Cancer Institute Clinical Proteomic Tumor Analysis Consortium (CPTAC), 2018), and The Cancer Genome Atlas from the Cancer Imaging Archive (TCGA-LGG & TGCA-GBM) (Pedano et al., 2016) (Scarpance et al., 2016).

## 4. Methods

### 4.1. BTReport Framework

BTReport is a two-stage approach for neuro-oncology RRG that first deterministically extracts interpretable and clinically relevant neuroimaging features from mpMRI, then uses an LLM for syntactically structured narrative construction and report generation. Since quantitative measurements are central to radiological evaluation, BTReport relies on validated open-sourced algorithms to derive descriptors of anatomy and tumor pathology, rather than performing end-to-end medical inference with a VLM.

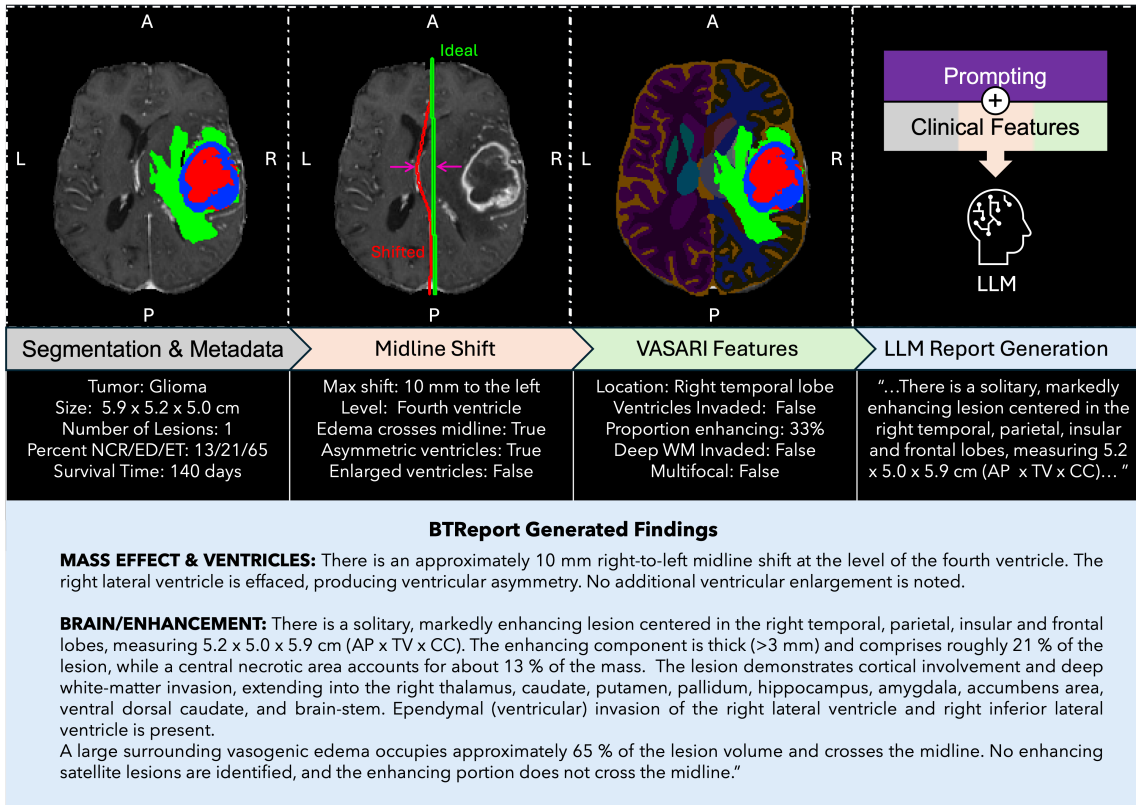


Figure 1: BTReport Overview: Interpretable, clinically meaningful variables are deterministically extracted for each case, including demographics, VASARI features, and 3D midline shift measurements. These features are utilized by context-guided LLMs for clinically grounded radiology report generation.

For a given case, the BTReport framework (Figure 1) consists of four steps:

1. T1-weighted scans and tumor masks are used to generate tumor-robust anatomical segmentations, from which regional volumetric statistics are computed (Section 4.2).
2. Midline shift is quantified by propagating a hand-annotated MNI152 midline into subject space using the MNI152-to-subject deformation field (Section 4.3).

3. Anatomical segmentations (Section 4.2) are used to extract VASARI (Visually AcceSable Rembrandt Images) features using a modified VASARI-auto pipeline (Ruffle et al., 2024), yielding standardized tumor- and anatomy-specific descriptors (Section 4.4).
4. A general-purpose LLM is provided a context-specific prompt and tasked with generating structured radiology reports, with generation conditioned on deterministic BTRreport-extracted features (Section 4.6).

## 4.2. Segmentation Statistics

**BTRreport Inputs:** For a given case, inputs to the BTRreport framework include their T1 sequence and corresponding tumor segmentation masks. Tumor segmentation masks are derived from manual annotation or automated segmentation, and tumor sub-region annotations (necrotic core, edema, enhancing tumor) follow BraTS convention.

**Robust Anatomical Segmentations & Volumetric Statistics:** Anatomical segmentations provide tumor-ROI relational context by localizing tumors with respect to critical brain structures. However, many automated methods assume near-normal anatomy and can fail when large tumors deform tissue and obscure boundaries. Prior work improves tumor-robust segmentations either by training with synthetic tumor-mask augmentation (Lei et al., 2024), or by synthesizing pseudo-healthy structural images that suppress tumor appearance, increasing the reliability of established segmentation pipelines (Iglesias et al., 2023). Our approach closely resembles the latter.

First, we register the MNI152 atlas (Collins et al., 1999; Fonov et al., 2011, 2009) to subject space using SynthMorph (Hoffmann et al., 2024). The resulting warped atlas, referred to as MNI152-to-subject, provides a pseudo-healthy anatomical representation of the brain. Next, we obtain anatomical segmentations by running SynthSeg (Billot et al., 2023) on the MNI152-to-subject volume. As SynthSeg operates on the tumor-free brain representations, anatomical segmentations remain robust when tumors deform native anatomy. Finally, anatomical labels are merged with the tumor and midline segmentations, to produce a unified mask that jointly represents normal neuroanatomy and pathological structures. With this joint segmentation, we can reliably extract quantitative features including tumor and ventricle volume, lesion count, tumor sub-region proportions, and tumor-ROI overlap.

## 4.3. 3D Midline Shift Measurement via Atlas-based Segmentation

Midline shift (MLS) is an intracranial pathology characterized by the displacement of brain tissue across the skull’s midsagittal axis. MLS arises as a result of traumatic brain injury or tumor mass effects and is an indirect indicator of elevated intracerebral pressure. Estimation of MLS is done by identifying the axial slice with the largest deviation, as indicated by midline structures such as the septum pellucidum, the third and fourth ventricles, and the falx cerebri. However, this estimation is subject to high inter-rater variability as there is not a standard procedure for axial slice level selection.

Here, we propose a novel pipeline for MLS estimation based on clinical guidelines, using a deep learning atlas-based segmentation approach. Our approach, shown in Figure 2, leverages the robust registration capabilities of SynthMorph (Hoffmann et al., 2024) to register hand-annotated midline segmentations from a MNI152 atlas template onto patient

T1 scans. These are compared to an “ideal” midline defined by connecting the anterior and posterior points of the falx cerebri for each axial slice. By calculating the distance between the ideal and subject midlines at each voxel, we obtain accurate 3D MLS estimations in seconds, giving a more complete picture in comparison to 2D automated or manual methods. This approach has strong zero-shot generalization and can be applied to MRI or CT scans.

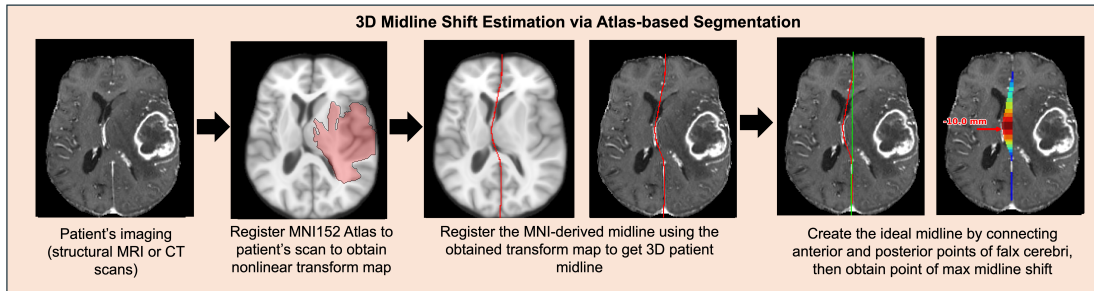


Figure 2: Atlas-based 3D midline shift (MLS) estimation using SynthMorph (Hoffmann et al., 2024), in which atlas midline annotations are registered to patient imaging and voxel-wise distances to an ideal midline axis are computed per axial slice.

#### 4.4. VASARI Feature Extraction

To standardize neuroimaging-derived glioma characterization and improve repeatability, the VASARI (Visually AcceSable Rembrandt Images) feature set was developed by The Cancer Imaging Archive. The VASARI feature set uses controlled vocabulary to quantitatively describes anatomical relationships between GBM tumors and clinically relevant brain structures established in the literature. Furthermore, these relationships are also routinely included in neuroradiology reports and used by neurosurgeons to assess whether patients are candidates for surgical intervention. VASARI features have been used to accurately predict tumor histological grade (WHO I-IV grade), disease progression, molecular mutation status (e.g., *IDH1* WT/mutant, *MGMT* un/methylated), risk of recurrence, and overall patient survival (Jain et al., 2014; Nicolasjilwan et al., 2015; Peeken et al., 2019; Setyawan et al., 2024; Wang et al., 2021; Zhou et al., 2017). Here, we employ a modified variant of VASARI-auto (Ruffle et al., 2024), an automated labeling tool, which has been validated as non-inferior to radiologist VASARI annotations, and can be used to reduce inter-rater variance. The included variations use the subject-space anatomical and midline segmentations extracted in Sections 4.2 and 4.3 to improve feature accuracy.

#### 4.5. BTReport Feature Validation

To determine whether extracted imaging features are clinically relevant, we assessed them based on two criteria: (1) whether features correspond to commonly reported radiological concepts (Section 4.5.1) and (2) by modeling their prognostic value with respect to patient survival (Section 4.5.2). To validate the accuracy of extracted features, we assess the quality of feature extraction in generated reports to those found in associated radiologist-authored ground truth reports, assessing categorical and numerical feature precision and error.

#### 4.5.1. SEMANTIC CLUSTERING OF COMMON RADIOLOGY FINDINGS

Here, we derived a ranked list of the most frequently included topics in real radiology reports. First, we extract all discrete factual claims across all reports in the HuskyBrain dataset using the method proposed in TBFact (Blondeel et al., 2025). TBFact uses an LLM (in this case DeepSeek-R1 (Guo et al., 2025)) to divide each clinical report into independently verifiable factual claims. Next, we embedded each extracted claim using a lightweight sentence-transformer model (all-MiniLM-L6-v2 (Reimers and Gurevych, 2019)) to obtain dense semantic representations. All embedded claims were pooled across the dataset and clustered using hierarchical agglomerative clustering (cosine distance, average linkage) without specifying the number of clusters a priori, allowing groups of semantically related statements to emerge from the data in an unsupervised manner.

We then use a pre-trained Gemma 3 27B LLM (Gemma Team et al., 2025) model to summarize claims in each cluster into representative topic sentences. For example, a cluster containing claims such as: [“The lesion measures 4.0 x 3.5 cm”, “The lesion measures 3.9 x 2.0 x 2.1 cm.”, ...], is summarized by the topic sentence “Lesion size and measurements reported.” The number of claims assigned to each cluster serves as an estimate of the frequency of the corresponding topic. The 35 most frequent cluster descriptors, along with representative example claims for each cluster, are reported in Appendix A.

Collectively, the resulting set of cluster descriptors and frequencies provides an interpretable, data-driven view of radiological concepts based on real-world reporting. We use these concepts to validate the clinical relevance of deterministically extracted features found in BTRReport and assess their alignment with findings in real radiology reports.

#### 4.5.2. SURVIVAL OUTCOME MODELING

To assess whether BTRReport-extracted features were predictive of clinical outcomes, we evaluated their association with overall survival measured from diagnosis. This analysis used the dataset described in Section 3.2, which provides survival outcomes and patient metadata for a subset of BraTS cases. Survival analyses were performed using Kaplan–Meier estimators implemented with the `lifelines` library (Davidson-Pilon, 2024), with curves reported in Appendix B. Differences between survival groups were assessed using the log-rank test, and relative risk was quantified using Cox proportional hazards models, reporting hazard ratios with 95% confidence intervals. Statistical significance was determined from p-values derived from both the log-rank tests and Cox models.

#### 4.5.3. CLINICALLY RELEVANT FEATURES INCLUDED IN BTRREPORT

Table 1 summarizes the quantitative imaging features used by BTRReport for automated report generation. Each feature was evaluated according to two clinically motivated criteria: 1) whether it corresponded to at least one of the top 35 radiological concepts most frequently documented in reference reports from the HuskyBrain dataset (Section 4.5.1), and 2) whether it was a statistically significant predictor of patient survival time based on the survival analysis described in Section 4.5.2. We found that 21/22 of the BTRReport features are commonly reported in clinical reports, and 11/22 are predictive of overall patient survival. This analysis highlights the clinical relevance of extracted features, motivating their use as inputs for downstream RRG.

Table 1: Summary of quantitative features used in BTReport. Feature groups are color-coded: gray = segmentation statistics, green = VASARI features, orange = midline features. Key: ● indicates the feature is among the top 35 most frequently reported concepts in real radiology reports, and ● indicates the feature is a significant survival predictor. Acronyms: WM-white matter; MLS-midline shift.

Total tumor volume (mL)	● ●	Ventricular Invasion	● ●	Level of max MLS	● ●
3D Lesion Sizes (cm)	●	Side of Tumor Epicenter	●	Max MLS (mm) + L/R	● ●
Proportion of Necrosis	●	Enhancement Quality	●	Edema crosses midline	● ●
Number of lesions	●	Enhancement thickness	●	ET Crosses midline	● ●
Proportion of Enhancing	● ●	Multiple satellites present	●	Asymmetrical Ventricles	●
Proportion of Edema	●	Multifocal or Multicentric	● ●	Enlarged Ventricles	● ●
Cortical involvement	●	Deep WM invasion	● ●		
Tumor Location	● ●	Eloquent Brain Involved	● ●		

#### 4.6. Structured Synthetic Report Generation Using LLMs

The BTReport framework generates the *Findings* sections of radiology reports in a style that mimics a target institution through an in-context learning prompt to provide stylistic guidance and emphasize the relevant radiological facts which should be included in generated reports (Appendix C). Some facts in reference reports (e.g., diffusion characteristics, vascular features) cannot be derived from standard structural mpMRI modalities used for tumor segmentation (T1w, T1c, T2w, T2-FLAIR). Therefore, the LLM is explicitly instructed to exclude these findings in generated reports to prevent unsupported claims.

Report generation is conditioned on the BTReport feature set and prompt instructions that restrict the LLM to narrative synthesis and stylistic alignment, enabling tighter control over report content and structure without task-specific fine-tuning. This prompting strategy improves interpretability by allowing reported findings to be directly traced to deterministic features.

For RRG with BTReport, we experiment with two open-source pre-trained LLMs with reasoning capability: gpt-oss:120b (Agarwal et al., 2025) and Llama 3.1 70B Instruct (Grattafiori et al., 2024). These models were selected based on their strong reasoning and instruction-following capabilities. To avoid cloud sharing of medical data, only local offline models were used. Examples *Findings* sections from synthetically generated reports across the different BTReport variants and other RRG frameworks can be found in Appendix F.

## 5. Evaluation

We evaluated BTReport on 129 paired GBM image-report cases from the HuskyBrain dataset. Synthetic reports generated by BTReport and other neuro-oncology RRG frameworks described in Section 2 were compared against ground-truth radiology reports using automated metrics (Section 5.1), and the reliability of BTReport-extracted features was assessed (Section 5.2). To test robustness to segmentation quality, we compared BTReport outputs generated using DeepMedic (Kamnitsas et al., 2017) segmentations, a dual-pathway CNN-based model, with outputs generated using segmentations from the BraTS 2023-winning submission (Ferreira et al., 2024), which is based on an nnU-Net (Isensee et al.,

2021) ensemble. Prior work (Pemberton et al., 2023) has shown nnU-Net-based models to outperform DeepMedic for brain tumor segmentation, motivating this comparison.

### 5.1. Automated Evaluation of Generated Radiology Reports

Generated radiology reports were evaluated using RadEval (Xu et al., 2025), a unified open-source framework that evaluates radiology text based on:

- N-gram-based lexical similarity: BLEU (Papineni et al., 2002), ROUGE (Lin, 2004)
- pre-trained contextualized embeddings: BERTScore (Zhang et al., 2020)
- clinically grounded scores: RaTEScore (Zhao et al., 2024)

To assess the clinical correctness of content in generated reports, we used TBFact (Blondeel et al., 2025), an LLM-based factuality metric that evaluates reports based on factual inclusion, distortion, and omission. These metrics are presented in Tables 3 and 4.

### 5.2. Reliability of Extracted Features

We evaluated the reliability of BTRReport-extracted clinical features by comparing them with the same features extracted from ground-truth radiology reports in the HuskyBrain dataset. Using LangExtract (Goel, 2025), we decompose radiology reports into class-attribute pairs, where each class represents a clinically defined radiology concept, and their associated attribute(s) specify the characteristic linked to the concept. For example, the text [“arising in the right parietal lobe...”] is decomposed into the “side of tumor epicenter” class and the “right” attribute (See Appendix E for a complete example). We present the accuracy of categorical features and the error of the numerical features in Table 2.

## 6. Results

### 6.1. Quality of Extracted Features

Table 2: Attribute-level reliability of BTRReport-extracted features relative to ground-truth HuskyBrain radiology reports. Categorical features report accuracy, and numeric features report mean absolute error (MAE)  $\pm$  standard deviation. Best-performing variant is highlighted in blue.

Method	Accuracy			MAE		
	Cortical involvement	Side of tumor epicenter	Ventricular effacement	Midline shift (mm)	Number of lesions	Tumor volume (cm <sup>3</sup> )
BTRReport (DM)	0.79	<b>1.00</b>	<b>0.56</b>	1.26 $\pm$ 2.60	1.48 $\pm$ 2.63	<b>1.80</b> $\pm$ 3.95
BTRReport (FI)	<b>0.80</b>	<b>1.00</b>	0.52	<b>1.21</b> $\pm$ 2.55	<b>0.41</b> $\pm$ 1.69	2.00 $\pm$ 5.16

DM: DeepMedic (Kamnitsas et al., 2017) segmentations    FI: Faking It (Ferreira et al., 2024) segmentations

Table 2 reports the reliability of BTRReport-extracted clinical features under two segmentation inputs. Across both variants, categorical features showed strong agreement with LangExtract-derived ground truth, including perfect accuracy for identifying the side of the tumor epicenter and high accuracy for cortical involvement and tumor location. Ventricular effacement was less reliable, indicating that this attribute may be more challenging to

describe consistently. For numeric features, BTReport achieved low mean absolute error for lesion count and midline shift, with consistently lower errors when using the improved segmentation model (FI). Tumor volume errors were comparable across variants.

## 6.2. Lexical Similarity of Generated Reports

Table 3: Mean  $\pm$  standard deviation BLEU and ROUGE metrics. Best values per metric are highlighted in blue. <sup>†</sup>Following Approximate Randomization, synthetic reports generated with all BTReport variants were superior to those generated using other frameworks across all evaluation metrics ( $p < 0.0001$ ).

Framework	BLEU-1	BLEU-2	ROUGE-1	ROUGE-2
BTReport (gpt-oss:120B, DM) <sup>†</sup>	0.236 $\pm$ 0.068	0.124 $\pm$ 0.044	0.360 $\pm$ 0.078	0.109 $\pm$ 0.038
BTReport (gpt-oss:120B, FI) <sup>†</sup>	0.244 $\pm$ 0.070	0.132 $\pm$ 0.044	<b>0.371 <math>\pm</math> 0.078</b>	<b>0.115 <math>\pm</math> 0.040</b>
BTReport (LLaMA3:70B, DM) <sup>†</sup>	<b>0.248 <math>\pm</math> 0.078</b>	<b>0.136 <math>\pm</math> 0.052</b>	0.362 $\pm$ 0.080	0.115 $\pm$ 0.043
AutoRG-Brain (Lei et al., 2024)	0.158 $\pm$ 0.072	0.080 $\pm$ 0.042	0.268 $\pm$ 0.060	0.070 $\pm$ 0.032
Seg-to-Exp (Valerio et al., 2025)	0.085 $\pm$ 0.039	0.035 $\pm$ 0.018	0.163 $\pm$ 0.055	0.023 $\pm$ 0.013

DM: DeepMedic (Kamnitsas et al., 2017) segmentations    FI: Faking It (Ferreira et al., 2024) segmentations

Table 3 summarizes the lexical similarity between generated and reference reports using mean BLEU and ROUGE metrics calculated over the 129 subject test dataset. Across all metrics, all BTReport variants substantially outperformed *AutoRG-Brain* and *Seg-to-Exp*, indicating closer n-gram overlap with clinical ground-truth reports. Across the two BTReport variants, BTReport with LLaMA3:70B achieved the highest BLEU-1 and BLEU-2 scores, indicating improved short-range lexical precision. In contrast, BTReport with gpt-oss:120B attained the highest ROUGE-1 and ROUGE-2 scores, suggesting improved recall of clinically relevant phrases and longer contextual spans. While the relative strengths marginally differed across metrics for BTReport variants, all demonstrated statistically significant improvements in lexical alignment with reference reports in comparison to baseline methods.

## 6.3. Factual Accuracy of Generated Reports

Table 4 reports factual consistency and semantic alignment of generated reports using TBFact, BERTScore, and RaTEScore. Across all metrics, both BTReport variants outperformed *AutoRG-Brain* and *Seg-to-Exp*, indicating improved factual grounding and clinically coherent report generation. BTReport with gpt-oss:120B and LLaMA3:70B achieved the highest overall performance, with the gpt-oss:120B model showing consistently higher TBFact, BERTScore, and RaTEScore values under the higher-quality segmentation input (FI) compared to DM. This suggests that reports generated with BTReport contain a greater proportion of verifiable clinical statements that are consistent with the underlying imaging findings. Both BTReport frameworks achieved substantially higher scores for the BERTScore and RaTEScore metrics, indicating closer semantic alignment with reference reports and improved clinical relevance. In contrast, *AutoRG-Brain* exhibited moderate

performance, while *Seg-to-Exp* showed limited factual consistency, reflected by low TBFact scores and reduced semantic similarity. All improvements using the BTRReport framework were statistically significant ( $p < 0.0001$ ) when evaluated with AR.

Table 4: Mean  $\pm$  standard deviation TBFact, BERTScore, and RaTEScore metrics. Best values per metric are highlighted in blue. <sup>†</sup> Following Approximate Randomization, synthetic reports generated with all BTRReport variants were superior to those generated using other frameworks across all evaluation metrics ( $p < 0.0001$ ).

Framework	TBFact (DeepSeek-R1)				BERTScore	RaTEScore
	Score	Prec.	Recall	F1		
BTRReport (gpt-oss:120B, DM) <sup>†</sup>	0.313 $\pm$ 0.145	0.345 $\pm$ 0.155	0.325 $\pm$ 0.163	0.313 $\pm$ 0.145	0.447 $\pm$ 0.060	0.568 $\pm$ 0.057
BTRReport (gpt-oss:120B, FI) <sup>†</sup>	<b>0.353</b> $\pm$ 0.151	<b>0.412</b> $\pm$ 0.146	<b>0.349</b> $\pm$ 0.162	<b>0.359</b> $\pm$ 0.145	<b>0.453</b> $\pm$ 0.055	<b>0.577</b> $\pm$ 0.054
BTRReport (LLaMA3:70B, DM) <sup>†</sup>	0.295 $\pm$ 0.130	0.377 $\pm$ 0.168	0.274 $\pm$ 0.136	0.295 $\pm$ 0.130	0.433 $\pm$ 0.062	0.567 $\pm$ 0.061
AutoRG-Brain (Lei et al., 2024)	0.072 $\pm$ 0.123	0.282 $\pm$ 0.155	0.186 $\pm$ 0.137	0.196 $\pm$ 0.130	0.327 $\pm$ 0.047	0.477 $\pm$ 0.053
Seg-to-Exp (Valerio et al., 2025)	0.014 $\pm$ 0.047	0.131 $\pm$ 0.121	0.147 $\pm$ 0.108	0.098 $\pm$ 0.089	0.156 $\pm$ 0.042	0.409 $\pm$ 0.039

DM: DeepMedic (Kamnitsas et al., 2017) segmentations    FI: Faking It (Ferreira et al., 2024) segmentations

## 7. Discussion

We present BTRReport, a framework for brain tumor RRG grounded in clinically relevant quantitative imaging features. Overall, our findings support the two-stage report generation paradigm for neuro-oncology RRG, suggesting that in medical imaging domains with limited data, adding quantitative features to prompts is an efficient way to generate reports and improve factual consistency. BTRReport enables accurate measurement-grounded reporting without requiring task-specific fine-tuning of vision encoders or VLMs, and generates interpretable reports by using modular, reliably-extracted features.

We presented novel pipelines for midline shift measurement and tumor-robust anatomical segmentation for deterministic feature extraction, and showed that these features had strong agreement with expert annotation. Additionally, other quantitative features were selected based on clinical guidelines such as VASARI and validated to be significant predictors of survival time using KM-analysis. By clustering concepts frequently reported in real radiology reports, we validated that the included features were clinically relevant. When compared with existing neuro-oncology RRG approaches, BTRReport generated reports are superior in terms of lexical similarity and factual accuracy. We believe this improvement is due to the clinically-relevant features used for generation and the in-context learning prompt which allows reasoning LLMs to predict the most relevant features for RRG.

To facilitate further research in neuro-oncology RRG, we provide BTRReport-BraTS, an open-source companion dataset containing anatomical descriptors, metadata, and BTRReport generated reports for mpMRI cases in the BraTS’23 dataset. To assess the clinical applicability of BTRReport, we have developed BTRReview, a survey tool for radiologist assessment of generated neuro-oncology radiology reports (Appendix H). Future work will obtain radiologist feedback with BTRReview, incorporate additional features such as white matter hyperintensities and basal cistern status, handle additional mpMRI imaging modalities, and include descriptions of ischemic or hemorrhagic stroke findings.

## Acknowledgments

The authors are grateful to Caitlin Neher, Leonardo Schettini, and Ankush Jindal for their helpful discussions. The authors are grateful for Dr. Kambiz Nael for his clinical perspectives for improving the evaluation platform. The authors are also grateful to James K. Ruffle, for making the VASARI-Auto implementation available.

## 8. Ethics

This study’s activities were approved by the Institutional Review Board at the University of Washington (STUDY00022466). This research is in accordance with the principles embodied in the Declaration of Helsinki.

## 9. Funding Statement

The work of Juampablo Heras Rivera was partially supported by the U.S. Department of Energy Computational Science Graduate Fellowship under Award Number DE-SC0024386.

## References

- Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman, Diogo Almeida, Janko Altmerschmidt, Sam Altman, Shyamal Anadkat, et al. GPT-4 technical report. *arXiv preprint arXiv:2303.08774*, 2023.
- Maruf Adewole, Jeffrey D Rudie, Anu Gbdamosi, Oluyemisi Toyobo, Confidence Raymond, Dong Zhang, Olubukola Omidiji, Rachel Akinola, Mohammad Abba Suwaid, Adaobi Emegoakor, et al. The brain tumor segmentation (BraTS) challenge 2023: Glioma segmentation in sub-saharan Africa patient population (BraTS-Africa). *ArXiv*, pages arXiv–2305, 2023.
- Sandhini Agarwal, Lama Ahmad, Jason Ai, Sam Altman, Andy Applebaum, Edwin Arbus, Rahul K Arora, Yu Bai, Bowen Baker, Haiming Bao, et al. gpt-oss-120b & gpt-oss-20b model card. *arXiv preprint arXiv:2508.10925*, 2025.
- S. Bakas, C. Sako, H. Akbari, M. Bilello, A. Sotiras, G. Shukla, J. D. Rudie, N. Flores Santamaria, A. Fathi Kazerooni, S. Pati, S. Rathore, E. Mamourian, S. M. Ha, W. Parker, J. Doshi, U. Baid, M. Bergman, Z. A. Binder, R. Verma, and C. Davatzikos. Multi-parametric magnetic resonance imaging (mpMRI) scans for de novo Glioblastoma (GBM) patients from the University of Pennsylvania Health System (UPENN-GBM). 2021. Dataset on The Cancer Imaging Archive (TCIA).
- Pedro RAS Bassi, Mehmet Can Yavuz, Ibrahim Ethem Hamamci, Sezgin Er, Xiaoxi Chen, Wenxuan Li, Bjoern Menze, Sergio Decherchi, Andrea Cavalli, Kang Wang, et al. RadGPT: Constructing 3d image-text tumor datasets. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 23720–23730, 2025.
- Benjamin Billot, Douglas N Greve, Oula Puonti, Axel Thielscher, Koen Van Leemput, Bruce Fischl, Adrian V Dalca, Juan Eugenio Iglesias, et al. SynthSeg: Segmentation

- of brain MRI scans of any contrast and resolution without retraining. *Medical image analysis*, 86:102789, 2023.
- Matthias Blondeel, Noel Codella, Sam Preston, Hao Qiu, Leonardo Schettini, Frank Tuan, Wen-wai Yim, Smitha Saligrama, Mert Öz, Shrey Jain, et al. Healthcare Agent Orchestrator (HAO) for Patient Summarization in Molecular Tumor Boards. *arXiv preprint arXiv:2509.06602*, 2025.
- E. Calabrese, J. Villanueva-Meyer, J. Rudie, A. Rauschecker, U. Baid, S. Bakas, S. Cha, J. Mongan, and C. Hess. The University of California San Francisco Preoperative Diffuse Glioma MRI (UCSF-PDGM). 2022. Dataset on The Cancer Imaging Archive (TCIA).
- Boyuan Chen, Zhuo Xu, Sean Kirmani, Brain Ichter, Dorsa Sadigh, Leonidas Guibas, and Fei Xia. Spatialvlm: Endowing vision-language models with spatial reasoning capabilities. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 14455–14465, 2024.
- Zhihong Chen, Yan Song, Tsung-Hui Chang, and Xiang Wan. Generating radiology reports via memory-driven transformer. *arXiv preprint arXiv:2010.16056*, 2020.
- Eric W Christensen, Jay R Parikh, Alexandra R Drake, Eric M Rubin, and Elizabeth Y Rula. Projected US radiologist supply, 2025 to 2055. *Journal of the American College of Radiology*, 22(2):161–169, 2025.
- D Louis Collins, Alex P Zijdenbos, Wim FC Baaré, and Alan C Evans. ANIMAL+ INSECT: improved cortical structure segmentation. In *Biennial International Conference on Information Processing in Medical Imaging*, pages 210–223. Springer, 1999.
- C. Davatzikos, S. Rathore, S. Bakas, S. Pati, M. Bergman, R. Kalarot, P. Sridharan, A. Gastounioti, N. Jahani, E. Cohen, H. Akbari, B. Tunc, J. Doshi, D. Parker, M. Hsieh, A. Sotiras, H. Li, Y. Ou, R. K. Doot, M. Bilello, Y. Fan, R. T. Shinohara, P. Yushkevich, R. Verma, and D. Kontos. Cancer imaging phenomics toolkit: quantitative imaging analytics for precision diagnostics and predictive modeling of clinical outcome. *Journal of Medical Imaging*, 5(1):011018, 2018. doi: 10.1117/1.JMI.5.1.011018.
- Cameron Davidson-Pilon. lifelines: Survival Analysis in Python. 2024. Zenodo software release.
- Maria Correia de Verdier, Rachit Saluja, Louis Gagnon, Dominic LaBella, Ujjwal Baid, Nourel Hoda Tahon, Martha Foltyn-Dumitru, Jikai Zhang, Maram Alafif, Saif Baig, et al. The 2024 brain tumor segmentation (brats) challenge: Glioma segmentation on post-treatment MRI. *arXiv preprint arXiv:2405.18368*, 2024.
- Dina Demner-Fushman, Marc D. Kohli, Marc B. Rosenman, Sonya E. Shooshan, Laritza Rodriguez, Sameer Antani, George R. Thoma, and Clement J. McDonald. Preparing a collection of radiology examinations for distribution and retrieval. *Journal of the American Medical Informatics Association*, 23(2):304–310, March 2016. doi: 10.1093/jamia/ocv080.

- André Ferreira, Naida Solak, Jianning Li, Philipp Dammann, Jens Kleesiek, Victor Alves, and Jan Egger. How we won brats 2023 adult glioma challenge? just faking it! enhanced synthetic data augmentation and model ensemble for brain tumour segmentation. *arXiv preprint arXiv:2402.17317*, 2024.
- Vladimir Fonov, Alan C Evans, Kelly Botteron, C Robert Almlı, Robert C McKinstry, D Louis Collins, Brain Development Cooperative Group, et al. Unbiased average age-appropriate atlases for pediatric studies. *NeuroImage*, 54(1):313–327, 2011.
- Vladimir S Fonov, Alan C Evans, Robert C McKinstry, CR Almlı, and DL Collins. Unbiased nonlinear average age-appropriate brain templates from birth to adulthood. *NeuroImage*, 47:S102, 2009.
- Gemma Team, Aishwarya Kamath, Johan Ferret, Shreya Pathak, Nino Vieillard, Ramona Merhej, Sarah Perrin, Tatiana Matejovicova, Alexandre Ramé, and Morgane Rivière. Gemma 3 technical report. *arXiv preprint arXiv:2503.19786*, 2025.
- A. Goel. LangExtract (1.0.9), 2025. URL <https://doi.org/10.5281/zenodo.17015090>.
- Aaron Grattafiori, Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Alex Vaughan, et al. The LLaMa 3 herd of models. *arXiv preprint arXiv:2407.21783*, 2024.
- Daya Guo, Dejian Yang, Haowei Zhang, Junxiao Song, Ruoyu Zhang, Runxin Xu, Qihao Zhu, Shirong Ma, Peiyi Wang, Xiao Bi, et al. Deepseek-r1: Incentivizing reasoning capability in llms via reinforcement learning. *arXiv preprint arXiv:2501.12948*, 2025.
- Malte Hoffmann, Andrew Hoopes, Douglas N Greve, Bruce Fischl, and Adrian V Dalca. Anatomy-aware and acquisition-agnostic joint registration with SynthMorph. *Imaging Neuroscience*, 2:1–33, 2024.
- Juan E Iglesias, Benjamin Billot, Yaël Balbastre, Colin Magdamo, Steven E Arnold, Sudeshna Das, Brian L Edlow, Daniel C Alexander, Polina Golland, and Bruce Fischl. Synthsr: A public ai tool to turn heterogeneous clinical brain scans into high-resolution t1-weighted images for 3d morphometry. *Science advances*, 9(5):eadd3607, 2023.
- Fabian Isensee, Paul F. Jaeger, Simon A. A. Kohl, Jens Petersen, and Klaus H. Maier-Hein. nnU-Net: A Self-Configuring Method for Deep Learning-Based Biomedical Image Segmentation. *Nature Methods*, 18(2):203–211, 2021. doi: 10.1038/s41592-020-01008-z.
- Rajan Jain, Laila M Poisson, David Gutman, Lisa Scarpace, Scott N Hwang, Chad A Holder, Max Wintermark, Arvind Rao, Rivka R Colen, Justin Kirby, et al. Outcome prediction in patients with glioblastoma by using imaging, clinical, and genomic biomarkers: focus on the nonenhancing component of the tumor. *Radiology*, 272(2):484–493, 2014.
- Alistair E. W. Johnson, Tom J. Pollard, Seth J. Berkowitz, Nathaniel R. Greenbaum, Matthew P. Lungren, Chih-Ying Deng, Roger G. Mark, and Steven Horng. MIMIC-CXR, a de-identified publicly available database of chest radiographs with free-text reports. *Scientific Data*, 6:317, 2019. doi: 10.1038/s41597-019-0322-0. URL <https://doi.org/10.1038/s41597-019-0322-0>.

- Konstantinos Kamnitsas, Christian Ledig, Virginia FJ Newcombe, Joanna P Simpson, Andrew D Kane, David K Menon, Daniel Rueckert, and Ben Glocker. Efficient multi-scale 3D CNN with fully connected CRF for accurate brain lesion segmentation. *Medical image analysis*, 36:61–78, 2017.
- Jiayu Lei, Xiaoman Zhang, Chaoyi Wu, Lisong Dai, Ya Zhang, Yanyong Zhang, Yanfeng Wang, Weidi Xie, and Yuehua Li. AutoRG-Brain: Grounded Report Generation for Brain MRI. *arXiv preprint arXiv:2407.16684*, 2024.
- Chin-Yew Lin. ROUGE: A Package for Automatic Evaluation of Summaries. In *Text Summarization Branches Out*, pages 74–81, Barcelona, Spain, July 2004. Association for Computational Linguistics. URL <https://aclanthology.org/W04-1013>.
- Guanxiong Liu, Tzu-Ming Harry Hsu, Matthew McDermott, Willie Boag, Wei-Hung Weng, Peter Szolovits, and Marzyeh Ghassemi. Clinically accurate chest x-ray report generation. In *Machine Learning for Healthcare Conference*, pages 249–269. PMLR, 2019.
- Bjoern H Menze, Andras Jakab, Stefan Bauer, Jayashree Kalpathy-Cramer, Keyvan Farahani, Justin Kirby, Yuliya Burren, Nicole Porz, Johannes Slotboom, Roland Wiest, et al. The multimodal brain tumor image segmentation benchmark (BRATS). *IEEE transactions on medical imaging*, 34(10):1993–2024, 2014.
- National Cancer Institute Clinical Proteomic Tumor Analysis Consortium (CPTAC). The Clinical Proteomic Tumor Analysis Consortium Glioblastoma Multiforme Collection (CPTAC-GBM). 2018. Dataset on The Cancer Imaging Archive (TCIA).
- Manal Nicolasjilwan, Ying Hu, Chunhua Yan, Daoud Meerzaman, Chad A Holder, David Gutman, Rajan Jain, Rivka Colen, Daniel L Rubin, Pascal O Zinn, et al. Addition of MR imaging features and genetic biomarkers strengthens glioblastoma survival prediction in TCGA patients. *Journal of Neuroradiology*, 42(4):212–221, 2015.
- Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. BLEU: a Method for Automatic Evaluation of Machine Translation. In *Proceedings of the 40th Annual Meeting of the Association for Computational Linguistics, July 6-12, 2002, Philadelphia, PA, USA*, pages 311–318. ACL, 2002. URL <https://aclanthology.org/P02-1040/>.
- S. Pati, A. Singh, S. Rathore, A. Gastounioti, M. Bergman, P. Ngo, S. M. Ha, D. Bounias, J. Minock, G. Murphy, H. Li, A. Bhattarai, A. Wolf, P. Sridaran, R. Kalarot, H. Akbari, A. Sotiras, S. P. Thakur, R. Verma, R. T. Shinohara, P. Yushkevich, Y. Fan, D. Kontos, C. Davatzikos, and S. Bakas. The Cancer Imaging Phenomics Toolkit (CaPTk): Technical Overview. In *BrainLes 2019. Lecture Notes in Computer Science*, volume 11993, pages 380–394. Springer, 2020. doi: 10.1007/978-3-030-46643-5\\_38.
- N. Pedano, A. E. Flanders, L. Scarpace, T. Mikkelsen, J. M. Eschbacher, B. Hermes, V. Sinneros, J. Barnholtz-Sloan, and Q. Ostrom. The Cancer Genome Atlas Low Grade Glioma Collection (TCGA-LGG). <https://doi.org/10.7937/K9/TCIA.2016.L4LTD3TK>, 2016. Dataset on The Cancer Imaging Archive (TCIA).

- Jan C Peeken, Tatyana Goldberg, Thomas Pyka, Michael Bernhofer, Benedikt Wiestler, Kerstin A Kessel, Pouya D Tafti, Fridtjof Nüsslin, Andreas E Braun, Claus Zimmer, et al. Combining multimodal imaging and treatment features improves machine learning-based prognostic assessment in patients with glioblastoma multiforme. *Cancer medicine*, 8(1):128–136, 2019.
- Hugh G Pemberton, Jiaming Wu, Ivar Kommers, Dominique MJ Müller, Yipeng Hu, Olivia Goodkin, Sjoerd B Vos, Sotirios Bisdas, Pierre A Robe, Hilko Ardon, et al. Multi-class glioma segmentation on real-world data with missing mri sequences: comparison of three deep learning algorithms. *Scientific reports*, 13(1):18911, 2023.
- Nils Reimers and Iryna Gurevych. Sentence-bert: Sentence embeddings using siamese bert-networks. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing*. Association for Computational Linguistics, 11 2019. URL <https://arxiv.org/abs/1908.10084>.
- Chris Rorden et al. NiVue: WebGL2 based medical image viewer. <https://github.com/niivue/niivue>, 2024. GitHub repository.
- James K Ruffle, Samia Mohinta, Kelly Pegoretti Baruteau, Rebekah Rajiah, Faith Lee, Sebastian Brandner, Parashkev Nachev, and Harpreet Hyare. VASARI-auto: Equitable, efficient, and economical featurisation of glioma MRI. *NeuroImage: Clinical*, 44:103668, 2024.
- L. Scarpance, T. Mikkelsen, S. Cha, S. Rao, S. Tekchandani, D. Gutman, J. H. Saltz, B. J. Erickson, N. Pedano, A. E. Flanders, J. Barnholtz-Sloan, Q. Ostrom, D. Barboriak, and L. J. Pierce. The Cancer Genome Atlas Glioblastoma Multiforme Collection (TCGA-GBM). <https://doi.org/10.7937/K9/TCIA.2016.RNYFUYE9>, 2016. Dataset on The Cancer Imaging Archive (TCIA).
- Andrew Sellergren, Sahar Kazemzadeh, Tiam Jaroensri, Atilla Kiraly, Madeleine Traverse, Timo Kohlberger, Shawn Xu, Fayaz Jamil, Cían Hughes, Charles Lau, et al. MedGemma technical report. *arXiv preprint arXiv:2507.05201*, 2025.
- Nurhuda Hendra Setyawan, Lina Choridah, Hanung Adi Nugroho, Rusdy Ghazali Malueka, and Ery Kus Dwianingsih. Beyond invasive biopsies: using VASARI MRI features to predict grade and molecular parameters in gliomas. *Cancer Imaging*, 24(1):3, 2024.
- Xiaoyu Shi, Rahul Kumar Jain, Yinhao Li, Ruibo Hou, Jingliang Cheng, Jie Bai, Guohua Zhao, Lanfen Lin, Rui Xu, and Yen-wei Chen. TextBraTS: Text-Guided Volumetric Brain Tumor Segmentation with Innovative Dataset Development and Fusion Module Exploration. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 638–648. Springer, 2025.
- Karan Singhal, Shekoofeh Azizi, Tao Tu, S Sara Mahdavi, Jason Wei, Hyung Won Chung, Nathan Scales, Ajay Tanwani, Heather Cole-Lewis, Stephen Pfohl, et al. Large language models encode clinical knowledge. *Nature*, 620(7972):172–180, 2023.

- Alberto G Valerio, Katya Trufanova, Salvatore de Benedictis, Gennaro Vessio, and Giovanna Castellano. From segmentation to explanation: Generating textual reports from MRI with LLMs. *Computer Methods and Programs in Biomedicine*, page 108922, 2025.
- Jing Wang, Xiaoping Yi, Yan Fu, Peipei Pang, Huihuang Deng, Haiyun Tang, Zaide Han, Haiping Li, Jilin Nie, Guanghui Gong, et al. Preoperative magnetic resonance imaging radiomics for predicting early recurrence of glioblastoma. *Frontiers in Oncology*, 11: 769188, 2021.
- David Wu, Fateme Nateghi Haredasht, Saloni Kumar Maharaj, Priyank Jain, Jessica Tran, Matthew Gwiazdon, Arjun Rustagi, Jenelle Jindal, Jacob M Koshy, Vinay Kadiyala, et al. First, do NOHARM: towards clinically safe large language models. *arXiv preprint arXiv:2512.01241*, 2025.
- Justin Xu, Xi Zhang, Javid Abderezaei, Julie Bauml, Roger Boodoo, Fatemeh Haghighi, Ali Ganjizadeh, Eric Brattain, Dave Van Veen, Zaiqiao Meng, et al. RadEval: A framework for radiology text evaluation. In *Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 546–557, 2025.
- Tianyi Zhang, Varsha Kishore, Felix Wu, Kilian Q. Weinberger, and Yoav Artzi. BERTScore: Evaluating Text Generation with BERT. In *8th International Conference on Learning Representations, ICLR 2020, Addis Ababa, Ethiopia, April 26-30, 2020*. OpenReview.net, 2020. URL <https://openreview.net/forum?id=SkeHuCVFDr>.
- Weike Zhao, Chaoyi Wu, Xiaoman Zhang, Ya Zhang, Yanfeng Wang, and Weidi Xie. RaTEScore: A Metric for Radiology Report Generation. In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pages 15004–15019, 2024.
- Hao Zhou, Martin Vallières, Harrison X Bai, Chang Su, Haiyun Tang, Derek Oldridge, Zishu Zhang, Bo Xiao, Weihua Liao, Yongguang Tao, et al. MRI features predict survival and molecular markers in diffuse lower-grade gliomas. *Neuro-oncology*, 19(6):862–870, 2017.

## Appendix A. Cluster Analysis of Common Topics in Reference Reports

Cluster Description	Frequency	Two Example Sentences in this cluster
Lateral ventricle effacement/asymmetry; possible...	50	"The lateral ventricles are symmetric.", "The lateral ventricles are symmetric"
Midline shift presence and magnitude...	27	"There is a 2 mm left midline shift", "There is a susceptibility millimeters rightward midline shift."
Cisterns appear normal, no obstruction.	19	"The basal cisterns are patent.", "The basal cisterns are patent"
Peritumoral/Vasogenic edema present. *(This...	18	"There is mild-to-moderate associated edema", "The first mass has associated surrounding vasogenic edema"
Lesion size and measurements reported.	16	"The lesion measures 4.0 x 3.5 cm", "The lesion measures 3.9 x 2.0 x 2.1 cm."
No acute intracranial hemorrhage or infarct.	15	"There is no acute infarct", "No acute infarct is seen"
Multiple intracranial masses present bilaterally...	15	"There is a large irregular enhancing mass centered in the right frontal lobe", "There is a right frontal lobe mass"
Multiple, enhancing intracranial lesions present.	14	"The lesion originates in the anterior paramedian left frontal lobe", "There is a lesion in the right frontal lobe"
White matter disease, likely nonspecific etiology.	14	"There are scattered deep and periventricular white matter T2/FLAIR hyperintensities", "There is mild subcortical and periventricular white matter T2 FLAIR abnormality"
No mass effect/midline shift. (Alternatively: No...	13	"There is no shift in the brain", "There is no shift of the brain structures"
Ventricular system: normal or prominent...	13	"The ventricles, sulci, and cisterns are normal", "The remaining ventricles, sulci, and cisterns are normal"
Restricted diffusion within the lesion(s).	9	"The solid components of the lesion demonstrate moderate diffusion restriction", "The lesion has peripheral areas of mild diffusion restriction"
Frontal/Temporal lobe FLAIR signal abnormality	9	"There is extensive T2/FLAIR signal abnormality in the right frontotemporal lobes", "There is surrounding FLAIR signal hyperintensity inferiorly going into the temporal lobe and posteriorly."
Herniation syndromes present on imaging. (or...	8	"There is suggestion of transtentorial herniation", "There is leftward subfalcine herniation"
Mass dimensions and size measurements.	8	"The mass measures 2.4 x 3.4 cm in axial cross-section", "The mass measures approximately 6.6 x 4.7 cm in transverse dimensions and 4.6 cm craniocaudally"
Uncal medialization, potentially impacting...	8	"There is right uncal medialization", "There is medialization of the right uncus"
No acute intracranial hemorrhage present....	7	"No parenchymal hemorrhage is present", "There is no associated hemorrhage"
Mass effect on lateral ventricles.	7	"The mass extends along the ependymal surface of the right lateral ventricle", "The mass extends into the posterior horn of the right lateral ventricle"
Edema predominantly affecting frontal & temporal...	6	"There is perilesional edema along the predominantly anterior aspect of the medial frontal lobes", "The vasogenic edema extends to the frontal lobe"
Mass size and measurements. (Alternatively: Lesion...	6	"The mass measures 40 x 53 mm", "The mass measures approximately 58 x 44 x 44 mm"
Corpus callosum lesion, midline crossing/spread.	6	"The lesion extends into the splenium of the corpus callosum, crossing midline to the right", "There is ependymal spread along the body of the corpus callosum"
Diffusion restriction presence/absence &...	6	"There is no associated restricted diffusion", "There are areas of internal diffusion restriction and susceptibility"
Frontal horn effacement & ventricular asymmetry.	6	"There is partial effacement of the right frontal horn", "There are areas of subtle ependymal enhancement in the bilateral frontal horns"
Hemorrhagic lesion with restricted diffusion.	5	"The lesion restricts diffusion and has intralesional hemorrhage", "The mass is T2 hyperintense, contains multiple foci of internal hemorrhage, and demonstrates mottled diffusion restriction consistent with hypercellularity and/or necrosis."
Cistern effacement suggests mass effect. (Or, more...	5	"There is effacement of the right crural cistern", "There is partial effacement of the basal cisterns"
Sulcal effacement, widespread cortical...	5	"There is sulcal effacement involving the right parietal, posterior temporal, and occipital lobes", "There is mild sulcal effacement of the left occipital lobe."
Mass shows restricted diffusion on imaging. (Or...	5	"There are patchy foci of restricted diffusion within the mass", "The second mass has diffusion restriction"
Corpus callosum mass/involvement. (or simply:...	5	"The mass extends into the right-sided genu of the corpus callosum", "The mass extends along the splenium of the corpus callosum"
Basal ganglia involvement with signal abnormality.	5	"The signal abnormality extends into the right basal ganglia, right thalamus, right cerebral peduncle, and right midbrain", "The hyperintensity involves the bilateral basal ganglia, with greater involvement on the left"
Ventricular size and morphology assessment.	4	"The third ventricle is near slitlike", "There is complete effacement of the third ventricle"
Midline shift at foramen of Monro	4	"There is a negative millimeters leftward midline shift at the level of the foramen of Monroe", "There is some millimeters of rightward midline shift at the level of the foramen of Monro"
No acute hydrocephalus present.	4	"There is no evidence of acute hydrocephalus", "There is no hydrocephalus"
Corpus callosum FLAIR edema/hyperintensity...	4	"The surrounding T2/FLAIR signal is similar and extends to the left splenium, septum pellucidum, and superior corpus callosum", "There is subtle patchy T2 FLAIR hyperintensity along the right body of the corpus callosum"
Peripheral mass enhancement characteristics. (Or...	4	"The first mass has peripheral enhancement with a nodular solid enhancing component", "The mass has irregular somewhat nodular peripheral enhancement"
Right cerebral peduncle lesion/mass effect.	4	"There is mass effect on the right cerebral peduncle", "The lesion has questionable extension into the posterior right cerebral peduncle"

Table 5: Example sentences associated with the top 35 radiological concept clusters ranked by their prevalence in reference radiology reports.

## Appendix B. Kaplan-Meier Survival Analysis of Extracted Features

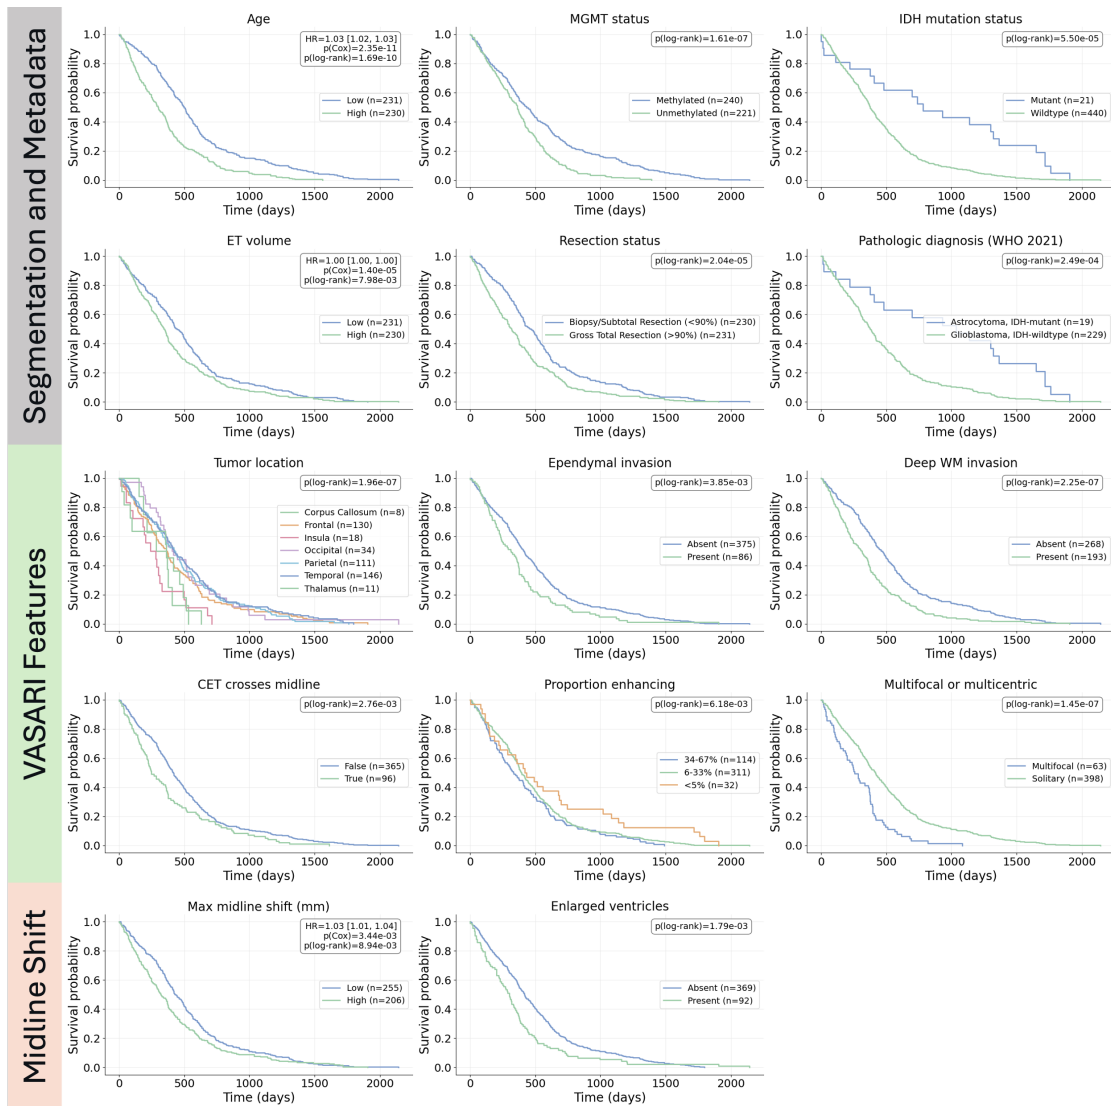


Figure 3: Kaplan-Meier (KM) plots for demographic and radiogenomic features, and deterministic features extracted with the BTRReport framework. Survival probabilities at specific time points are obtained by projecting vertically from the time of interest to the curve and horizontally to the y-axis. Example: for the Age feature, at 500 days post-diagnosis, the estimated survival probability for the older age group (High) is  $\sim 20\%$ , compared to  $\sim 50\%$  for the younger age group (Low). Kaplan-Meier analyses indicate that many of these features are predictive of overall survival, highlighting their clinical relevance and motivating their use as structured inputs for radiology report generation.

## Appendix C. LLM Prompt for Radiology Report Generation

**BTReport Prompt**

You are a radiologist generating a synthetic clinical MRI report.  
Below are example FINDINGS sections taken from real brain tumor reports:

**EXAMPLE FINDINGS:**  
{example\_findings}

Your job is to generate a FINDINGS section in the same clinical style, but only using the METADATA provided below.

Please abide strictly to the following rules (follow them exactly).

1. Use only the metadata provided for quantitative statements. Do NOT hallucinate any information that is not directly inferable.
2. Include 10–20 clinically meaningful findings summarized in an anatomically descriptive manner. Prioritize describing abnormal or clinically significant observations.
3. Preserve the subsection structure from the example reports. Make sure to include the following subsections: MASS EFFECT & VENTRICLES and BRAIN / ENHANCEMENT.
4. Never mention imaging sequences other than T1n, T2w, T2 FLAIR, or T1-Gd. Do not mention diffusion, perfusion, spectroscopy, MRA, or other modalities (unless stated explicitly in metadata).
5. Do not mention structures or measurements not present in the metadata.
6. Mandatory Considerations: Make sure to include the following findings below if present in metadata. Remember to follow the sentence structure in the example reports.
  - a) Maximum midline shift represented in mm units.
    - i) Make sure to describe the magnitude and the direction of the midline shift.
    - ii) Describe the anatomical level at which the (e.g., foramen of Monro, third and fourth ventricles, septum pellucidum).
    - iii) If shift is minimal (e.g., < 5 mm), explicitly state the measurement as no shift, but still provide the measurement.
  - b) If tumor mass effect is present, describe the mass effect on ventricles or surrounding brain structures.
    - i) Include a description of ventricular effacement (if present), including which horn (anterior/posterior horn), and on which hemisphere it is observed.
  - c) Comment on ventricular status. If effacement is present, describe the extent of the effacement or asymmetry. If ventricles are normal, explicitly state so (mirroring example reports). Use the following metadata fields in your description: “Asymmetrical Ventricles”, “Enlarged Ventricles”.
  - d) Describe the size of the primary lesion, as well as any smaller secondary lesions (if present) represented in cm units. Use the 3D measurements from the metadata. Make sure to include the following:
    - i) If multiple lesions exist, summarize number, dominant lesion, and laterality. Use the following metadata fields in your description: “Number of lesions” and “Multifocal or multicentric.”
    - ii) Anatomical location of lesion(s).
    - iii) Use the following metadata fields in your description: “Tumor Location”, “Side of Tumor Epicenter”, and “Region Proportions.”
  - e) Describe enhancing characteristics. Use the following metadata fields in your description: “Enhancement quality”, “Thickness of enhancing margin”, “Proportion Enhancing”.
    - i) Describe enhancement style (e.g., rim-enhancing, mildly enhancing, peripheral enhancing, multi-lobular enhancing) only if explicitly supported.
    - ii) Describe edematous tissue. Use the following metadata fields in your description: “ED volume”, “Whether edema crosses midline”, “Proportion of edema”.
    - iii) Describe vasogenic edema and its extent only if metadata supports it.
  - f) Describe invasion and involvement. Use the following metadata fields in your description: “Cortical involvement”, “Deep WM invasion”, “Ependymal invasion”, “Eloquent Brain Involvement”.
  - g) Describe the necrosis if present. Use the metadata field “Proportion Necrosis” to describe the central foci of necrosis.

**METADATA (for subject {subject\_id}):**  
{metadata\_json}

Write the **FINDINGS** section now, using clinical radiology language.

## Appendix D. Short version of LLM Prompt for BTReport

### BTReport Prompt (Short)

You are a radiologist generating a synthetic clinical MRI report.  
Below are example FINDINGS sections taken from real brain tumor reports:

**EXAMPLE FINDINGS:**  
{example\_findings}

Now generate a similar FINDINGS section, but **only** using the metadata provided below.

- Do not hallucinate any information that is not directly inferable.
- Preserve all subsections present in the example findings reports.
- Select the top 10 metadata-supported findings; real reports typically include 7–10 facts.
- Prioritize abnormal or clinically significant findings.
- Only T1n, T2w, T2 FLAIR, and T1-Gd sequences were obtained; do not comment on diffusion or other modalities.
- Comment on midline shift in the style of the example reports, including direction and magnitude.
- Describe mass effect and ventricular effacement if present, including laterality and anterior/posterior horn involvement.
- Report lesion dimensions in three orthogonal axes (cm) if supported by metadata.
- Do not mention structures, measurements, or features unless supported by the metadata.

**METADATA (for subject {subject\_id}):**  
{metadata\_json}

Write the **FINDINGS** section now, using clinical radiology language.

## Appendix E. Radiology report feature extraction for validating BTReport features

Highlights Legend: deep\_white\_matter\_invasion midline\_shift multifocality multiple\_satellite\_lesions number\_of\_lesions side\_of\_tumor\_epicenter tumor\_dimensions tumor\_location ventricle\_symmetry ventricular\_effacement

**BRAIN:**  
Somewhat ill-defined enhancing left frontal parietal mass lesion measuring 4.0 x 3.0 x 3.6 cm (APxTVsCC). Surrounding nonenhancing T2/FLAIR white matter hyperintensity. Speckled areas of susceptibility signal dropout correspond to areas of calcification on same-day noncontrast head CT. Surrounding neovascularity is present.

No additional intracranial lesions. No acute infarct.

**MASS EFFECT & VENTRICLES:**  
3 mm rightward midline shift. No transaxial herniation. Left-sided sulcal effacement. Partial effacement of the left lateral ventricle.

Figure 4: Visualization of features extracted from free text radiologist-authored report using LangExtract.

Highlights Legend: cortical\_involvement deep\_white\_matter\_invasion midline\_shift multifocality multiple\_satellite\_lesions number\_of\_lesions side\_of\_tumor\_epicenter tumor\_dimensions tumor\_location ventricle\_enlargement ventricle\_symmetry ventricular\_effacement ventricular\_invasion

**BRAIN:**  
**/ENHANCEMENT:**  
A solitary, left-sided lesion involves the frontal, parietal, and cortical regions. The mass measures 4.1 x 3.5 x 3.8 cm (AP x TV x CC). Marked enhancement is present with a thick (>3 mm) enhancing margin. A small central component of necrosis is noted. The lesion demonstrates deep white-matter invasion as well as cortical involvement. There is no ependymal (ventricular) invasion and the enhancing component does not cross the midline. Surrounding vasogenic edema is extensive, comprising the majority of the abnormal signal, but does not cross the midline. No additional satellite lesions are identified.

**MASS EFFECT & VENTRICLES:**  
There is a minimal left-to-right midline shift measuring approximately 4 mm at the level of the septum pellucidum. The left lateral ventricle is effaced, producing asymmetry of the ventricular system. No tonsillar herniation is identified and the basal cisterns remain patent.

Figure 5: Visualization of features extracted from free text BTReport-generated report using LangExtract.

## Appendix F. Example Generated Reports

### Model comparison of RRG output: Findings Section

**COLOR KEY:** Green indicates correct statements; yellow partially correct; red incorrect; gray ground-truth statements referenced by generated reports; **\*\*** findings derived from unavailable MRI modalities.

#### Ground Truth Findings

**MASS EFFECT & VENTRICLES:** Prominent leftward midline shift by approximately 14 mm (XXX/XXX). There is also medialization of right uncus. The basal cisterns are partially effaced.

**BRAIN/ENHANCEMENT:** A large irregular enhancing lesion centered within the right temporal lobe with significant mass effect. Restricted diffusion\*\* is noted within the enhancing portion of the lesion, which contains susceptibility artifact suggestive of microhemorrhages/angioinvasion.\*\* The lesion measures approximately 6.8 x 4.4 x 4.8 cm (AP, TV, CC). There is probable small subependymal enhancement at the atria of right lateral ventricle (XXX/XXX, XXX/XXX). No acute hematoma or infarct is seen.

#### Generated Findings

Framework	Prediction
<b>BTRreport</b> (gpt-oss:120B)	ROUGE: 0.37 BLEU: 0.24 RATESCR: 0.56 BERTSCR: 0.47 TB-F1: 0.33 <b>MASS EFFECT &amp; VENTRICLES:</b> Approximately <span style="background-color: #fff2cc;">12 mm</span> of <span style="background-color: #d9ead3;">right-to-left midline shift</span> is present, measured at the level of the falx cerebri above. The ventricles are asymmetrical, with compression of the right lateral ventricle secondary to <span style="background-color: #d9ead3;">tumor-related ependymal invasion</span> ; there is no overall ventricular enlargement. No tonsillar herniation is identified. <b>BRAIN/ENHANCEMENT:</b> A solitary <span style="background-color: #d9ead3;">right-sided lesion</span> involving the <span style="background-color: #d9ead3;">temporal, cortical, and parietal lobes</span> measures <span style="background-color: #fff2cc;">7.1 x 5.6 x 5.3 cm (AP x TV x CC)</span> . The mass shows marked <span style="background-color: #d9ead3;">heterogeneous enhancement</span> with a thick (> 3mm) enhancing margin; roughly 28% of the tumor volume enhances. A necrotic core comprises about 22% of the lesion volume. Deep white-matter invasion and cortical involvement are evident, and there is <span style="background-color: #d9ead3;">direct ependymal invasion of the adjacent lateral ventricle</span> . The enhancing component does not cross the midline, but extensive vasogenic edema (84 mL, around 50% of total lesion volume) crosses the midline and surrounds the lesion, extending into the contralateral hemisphere. Multiple small enhancing satellite nodules are present along the cortical surface. The edema involves visual-association cortex, indicating eloquent brain (vision) involvement.

**BTRreport  
(LLaMa3:70B)**

ROUGE: 0.42 BLEU: 0.29 RATESCR: 0.57 BERTSCR: 0.47 TB-F1: 0.34

**MASS EFFECT & VENTRICLES:** There is a maximum of 12 mm leftward midline shift at the level of the falx cerebri above. The ventricles are asymmetrical, with tumor-side ventricular asymmetry noted. There is no enlargement of the ventricles.

**BRAIN/ENHANCEMENT:** Within the right temporal, cortex, and parietal lobes, there is a solitary, predominantly enhancing mass measuring 7.1 x 5.6 x 5.3 cm. The lesion exhibits marked enhancement with a thick (> 3mm) enhancing margin. There are multiple small satellite lesions present. The enhancing region does not cross the midline. The lesion invades the cortex and deep white matter, with ependymal invasion also present. There is a large volume of surrounding vasogenic edema (83.6 mL), which crosses the midline. The proportion of edema is 49.74%. A central focus of necrosis is present, comprising approximately 21.99% of the lesion.”

**AutoRG-Brain**

ROUGE: 0.32 BLEU: 0.25 RATESCR: 0.53 BERTSCR: 0.41 TB-F1: 0.25

After contrast administration, the lesion shows significant enhancement on T1-weighted images with unclear boundaries. The adjacent brain parenchyma shows significant edema in the brain pools and ventricles. After contrast administration, the lesion in the right temporal lobe demonstrates significant ring-like enhancement. The lesion measures approximately 82\*111\*82mm, with indistinct borders and a large amount of edema in the adjacent brain tissue, causing compression of the right lateral ventricle and deviation of midline structures to the left.

**Seg-to-Exp**

ROUGE: 0.21 BLEU: 0.16 RATESCR: 0.41 BERTSCR: 0.26 TB-F1: 0.0

The brain tumor segmentation results indicate that the tumor is primarily located in the right temporal and parietal regions, with a significant presence in the Temporal-to-Parietal (GapMap) region (40.03% of the tumor). The tumor also affects several surrounding regions, including Area STS2 (STS), Area STS1 (STS), Frontal-to-Temporal-II (GapMap), and Area TE 2.2 (STG).

Semantic segmentation of the tumor reveals three distinct categories: **Tumor Core** (red), **Peritumoral Edema** (yellow), and **GD-Enhancing Tumor** (green). The Tumor Core represents the most aggressive part of the tumor, while Peritumoral Edema indicates the area surrounding the tumor that is affected by edema. The GD-Enhancing Tumor region shows enhancement on Gadolinium (GD) MRI sequences, suggesting a highly vascularized and active part of the tumor.

**Appendix G. BTRreport-BraTS: A Companion Dataset for BraTS RRG**

Pre-operative mpMRI cases from the combined training and validation splits of the BraTS 2023 Adult Glioma (BraTS'23) dataset (n=1,470 cases) were used to develop the BTRreport-BraTS dataset, an open-source companion dataset generated using the BTRreport framework. For each case, corresponding midline segmentations, extracted metadata, structured summary reports, and radiology reports generated using BTRreport are openly available in the project GitHub page.

## Appendix H. BTReview - A Tool for Radiologist Assessment of Generated Brain Tumor Radiology Reports

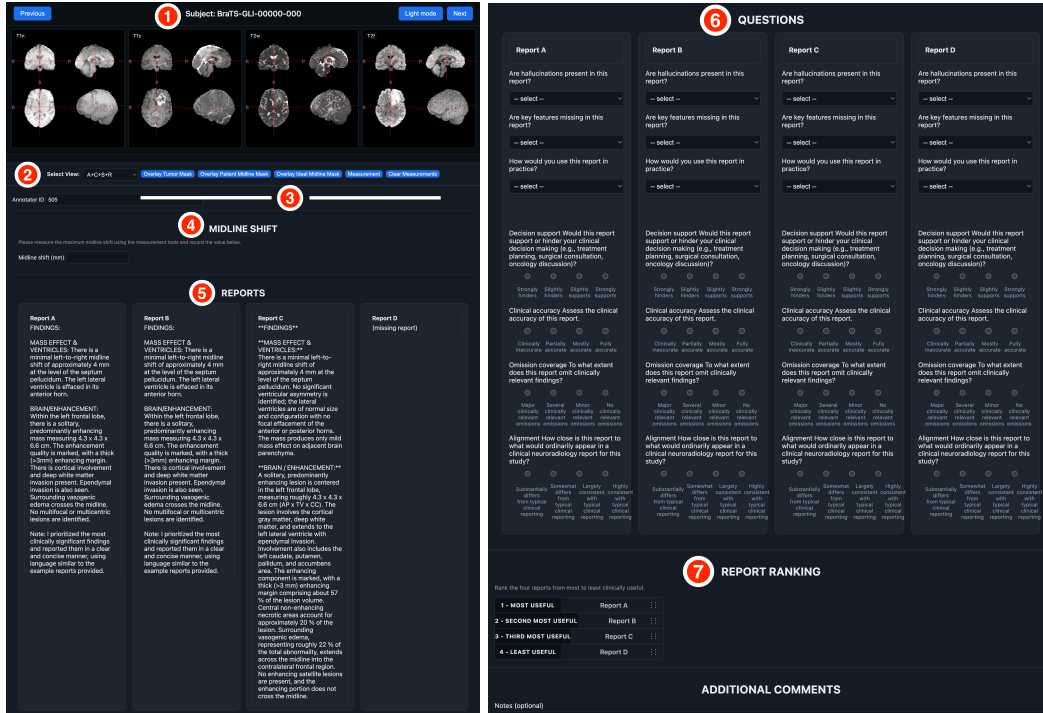


Figure 6: A clinical assessment tool to evaluate the clinical quality of synthetically generated radiology reports generated using various RRG frameworks.

Despite rapid advances in radiology report generation (RRG), there remains a lack of accessible, standardized tools that enable radiologists to efficiently review, compare, and score synthetically generated reports across different RRG frameworks. To address this gap, we developed BTReview, an open-source, web-based platform that supports de-identified case review and double-blinded clinical assessment of generated reports. This platform is customizable with any set of glioblastoma mpMRI cases and is designed to record structured, high-quality radiologist evaluations assessing the overall clinical quality of radiology reports generated with any RRG framework. The generation of these radiologist assessments creates a practical pathway for producing clinically grounded data and error annotation at scale, essential for generating high quality data for VLM training. Furthermore, this validation approach critically evaluates the translatability of existing RRG frameworks to ensure alignment with real-world clinical scenarios. We denote key features of BTReview platform with numerical labels and describe these features below. Feedback from a board-certified neuroradiologist was incorporated into BTReview during the design stage to ensure clinical relevance, appropriate terminology, and alignment with real-world neuroradiology reporting practices.

1) **Interactive Multi-View:** Leveraging the capabilities of NiiVue (Rorden et al., 2024), the first panel is an interactive multi-view of T1n, T1c, T2w, and T2f image sequences that can run on any device (phone, tablet, computer), enabling radiologists to review de-identified mpMRI cases in various clinical environments.

2) **View Selector:** Users can select and modify their multi-view from the drop down menu, enabling them to view image sequences for a given case in the axial, coronal, and sagittal planes, as well as in a 3D render view. An additional option allows all planes to be viewed simultaneously for a given case. Changes in view selection will propagate across all image sequences.

3) **Clinical Toolbar:** Within the clinical toolbar, radiologists can select blue buttons to overlay tumor masks, BTRReport-generated midline, and the brain’s ideal midline on image sequences in the interactive multi-view. Users can toggle the button again to remove the overlays. Additionally, a measurement tool enables clinicians to obtain quantitative measurements, such as midline shift by dragging their cursor between two points on any image sequence. These measurements will be automatically updated in the multi-view pane, allowing radiologists to revisit measurements across previously annotated cases. Users can also clear their measurements and perform re-measurements if desired.

4) **Midline Shift Measurement:** To validate midline shift measurements generated with BTRReport, radiologists can overlay the tumor mask and the brain’s ideal midline and scroll to a slice that demonstrates strong midline deviation (e.g., max midline shift). Users can then use the measurement tool to directly obtain a quantitative distance measurement from the boundaries of different sub-regions of the tumor to boundaries demarcated by midline brain structures (e.g., septum pellucidum) for midline shift measurement. Different sub-regions (necrotic core, edema, and enhancing tumor sub-regions) of the tumor mask are color coded and derived from BraTS tumor segmentation. Note: the appearance of tumor segmentation tasks may appear differently (e.g., masks from some institutions may not provide granular parcellation of tumor sub-regions). Radiologists can then overlay the BTRReport-generated midline and compare the magnitude of midline deviation through direct measurement and compare their midline shift measurement to the measurement described in synthetically generated radiology reports (see below).

5) **RRG Comparison:** To ensure unbiased evaluation of the clinical quality of synthetically generated radiology reports, the *Findings* description of the reports generated by the BTRReport(gpt-oss:120B), BTRReport(Llama 3:70B), AutoRG-Brain, and Seg-to-Exp RRG pipelines are displayed, where each column represents the synthetic output generated by a different pipeline. Importantly, synthetically generated reports are described as Report A, B, C, etc., ensuring reviewers are blinded to the identity of the pipeline used to generate the displayed reports. To mitigate potential bias during the review stage, the order in which the radiology report are shown (from left to right) is randomized for each case, ensuring that reviewers are not entrained to any associations with the generated reports and the order in which the reports are displayed.

6) **Clinical Assessment:** After radiologists have had the chance to review image sequences, they will proceed with the clinical assessment, where they will review radiology reports generated through different RRG frameworks. After their review, radiologists will then address a series of questions designed to critique the clinical quality of each generated report based on the following:

a) Generally, the first three questions aim to assess the real-world limitations of synthetic reports generated using automated RRG, and whether radiologists would consider reports clinically useful.

The first question will ask reviewers whether they identified hallucinations across the generated reports. If the radiologist selects “minor” or “major”, a conditional question will appear, prompting radiologists to select the types of hallucinations they observed. The next question asks whether key radiological features are missing in this report. Similar to the previous question, if the reviewer selects “some” or “many”, a follow up question will appear, prompting them to select from a list of missing elements that would be impactful for improving report quality. The final question asks reviewers to consider how they would personally use each generated report in clinical practice. Here, they will have the opportunity to select from the following responses from the drop down menu: “As a first draft”, “As a cross check / second reader”, “As a summary aid”, and “Would not use.”

Summarized below are the questions and their conditional question response options.

Question	Response Options (single-select)	Conditional Follow-Up (select all that apply)
<b>Hallucinations</b>	None; Minor; Major	<p><i>Shown if Minor/Major selected:</i></p> <p><b>Type of hallucination(s) observed?</b></p> <ul style="list-style-type: none"> <li>• Incorrect anatomical location of tumor</li> <li>• Incorrect tumor characteristics (e.g., size, laterality, enhancement)</li> <li>• Incorrect clinical implication</li> <li>• Fabricated finding</li> <li>• Other (free text: <i>Other hallucination details</i>)</li> </ul>
<b>Missing Features</b>	No; Some; Many	<p><i>Shown if Some/Many selected:</i></p> <p><b>Most impactful missing element(s)?</b></p> <ul style="list-style-type: none"> <li>• Tumor size/extent</li> <li>• Enhancement characteristics</li> <li>• Edema/mass effect</li> <li>• Midline shift</li> <li>• Multifocality</li> <li>• Invasion/eloquent cortex</li> <li>• Other (<i>free text</i>)</li> </ul>
<b>Intended use</b>	As a first draft; As a cross-check/second reader; As a summary aid only; Would not use	—

Table 7: Clinical evaluation questions and response options, including conditional follow-up items triggered by reviewer selection.

b) Next, radiologists will complete a series of Likert-scale questions asking them to evaluate each generated report based on the following criteria: (a) decision support, (b) clinical accuracy, (c) omission coverage, and (d) clinical structure.

**Decision Support:** Reviewers will be asked to evaluate whether the *Findings* narrative would positively or negatively affect downstream clinical decision making processes (e.g., treatment planning, surgical consult, oncology evaluation).

**Clinical Accuracy:** Reviewers will evaluate the factual accuracy of generated reports.

**Omission Coverage:** Reviewers will evaluate whether clinically relevant findings important for overall radiological interpretation of tumor effects on the surrounding brain environment are included in generated reports, or whether critical clinical findings are lacking or missing.

**Clinical Structure:** The last question asks reviewers to assess the extent to which generated *Findings* narrative differ or align with the clinical organization and conventions typically used in reports.

Summarized below are the question response options for their associated Likert-scale question.

Question	Criteria	1	2	3	4
Would this report support or hinder your clinical decision making (e.g., treatment planning, surgical consultation, oncology consultation)?	<b>Decision Support</b>	Strongly hinders	Slightly hinders	Slightly supports	Strongly supports
Assess the clinical accuracy of this report.	<b>Clinical Accuracy</b>	Clinically inaccurate	Partially inaccurate	Mostly accurate	Fully accurate
To what extent does this report omit clinically relevant findings?	<b>Clinical Omission</b>	Major clinically relevant omissions	Several clinically relevant omissions	Minor clinically relevant omissions	No clinically relevant omissions
To what extent does this report differ or align with the clinical structure of a standard neuroradiology report?	<b>Clinical Structure</b>	Substantially differs from typical clinical reporting	Somewhat differs from typical clinical reporting	Somewhat consistent with typical clinical reporting	Highly consistent with typical clinical reporting

Table 8: Four-point Likert-scale response options used for radiology report clinical assessment.

7) **Report Ranking:** As a final part of the clinical assessment, radiologists will be asked to rank generated reports based on clinical usefulness. In this question, they will be asked to order reports by “Most useful”, “Second most useful”, “Third most useful”, and “Least useful”. At the end of their clinical assessment, radiologists will be given the opportunity to provide optional comments and notes for the case they reviewed before saving their annotation and proceeding to the next case for review. The platform is designed with system memory, meaning that reviewers can revisit previously reviewed cases and make changes to their clinical assessments and annotations. All recorded assessments will be saved as a .json file when the browser closes.