

GLI-AL: A Multi-Modal Glioma MRI Label Resource with Unified Anatomy-Lesion Labels

Xingyu Xiang^{2,3}, Shuang Hao^{1,3}, Fan Wang^{1,3}, Jianhua Ma^{1,3}, Chunfeng Lian^{2,3}

1 Key Laboratory of Biomedical Information Engineering of Ministry of Education, School of Life Science and Technology, Xi'an Jiaotong University, Xi'an, China

2 School of Mathematics and Statistics, Xi'an Jiaotong University, Xi'an, China

3 Research Center for Intelligent Medical Equipment and Devices (IMED), Xi'an Jiaotong University, Xi'an 710049, China

Abstract

Existing BraTS-GLI datasets provide a widely used benchmark for adult glioma MRI segmentation, but their task definition focuses on tumor subregions and does not systematically represent coexisting white matter hyperintensities (WMH). In joint segmentation settings, such unlabeled abnormalities introduce task-specific label noise by treating pathological regions as normal tissue. To address this limitation, we introduce BraTS-GLI Anatomy-Lesion, a controlled-access, labels-only derived resource built from the BraTS 2023-GLI training cohort. The resource provides 1,251 unified eight-class anatomy-lesion label sets aligned with the original four-modal MRI cases, including image-repair labels for 116 cases requiring repaired imaging inputs. The cohort is organized into a 394-case purified subset and an 857-case extended subset, with case-level metadata covering label source, image-repair requirements, quality-control status, access conditions, checksums, and release boundaries. Compared with the original BraTS-GLI annotations, the resource substantially expands foreground supervision by incorporating healthy brain tissues and previously unlabeled coexisting abnormalities within a unified label space. A validation study using MedNeXt and T1/FLAIR inputs suggests that WMH-aware supervision preserves healthy-tissue segmentation performance across both in-domain GLI and external WMH datasets, while improving sensitivity to coexisting lesions relative to noisy-control training. The resource is intended for scientific research and supports joint anatomy-lesion supervision, label-noise analysis, and reproducible evaluation. Data are available at <https://www.synapse.org/Synapse:syn75210889/wiki/>, and code is available at <https://github.com/xyx200/brats-gli-anatomy-lesion-code>. The data resource DOI is <https://doi.org/10.7303/SYN75210889>.

Keywords

Anatomy-Lesion Segmentation, Label Noise, BraTS-GLI, White Matter Hyperintensities

Article informations

<https://doi.org/10.59275/j.melba.2026-575f> ©2026 Xingyu Xiang, Shuang Hao, Fan Wang, Jianhua Ma, and Chunfeng Lian. License: CC-BY 4.0

Volume 2026, Received: 2026-06-19, Published 2026-09-21

Corresponding author: chunfeng.lian@xjtu.edu.cn

Special issue: MICCAI Open Data 2026 x MELBA

Guest editors: AT, MS et al.



1. Background

BraTS datasets provide multi-center, pre-operative, multi-parametric MRI and expert tumor-subregion annotations for brain tumor segmentation research, and they are among the central public benchmarks for machine learning in glioma imaging (Menze et al., 2015; Bakas et al., 2017a,b,c; Baid et al., 2021; Karargyris et al., 2023). BraTS 2023-GLI contains four-modal images of adult glioma cases, with

tumor annotations covering necrotic tumor core (NCR), peritumoral edema (ED), and enhancing tumor (ET) regions (Baid et al., 2021). This task definition is highly effective for glioma segmentation, but it also raises an issue that is easily overlooked in joint segmentation settings: datasets usually annotate only the target lesion and do not systematically annotate other abnormalities that may coexist in the patient's brain (Rudie et al., 2019; Baid et al., 2021).

Table 1: Resource-context comparison

Resource	Anatomy labels	Primary target lesion labels	Comorbid lesion labels	Case stratification
FreeSurfer/SynthSeg	✓	×	×	×
BraTS	×	✓	×	×
SOOP stroke dataset	×	✓	×	×
GLI-AL	✓	✓	✓	✓

WMH is a common coexisting abnormality in BraTS-GLI brain MRI. Rudie et al. added expert WMH annotations to 285 BraTS 2018 training cases and reported that 68.8% of the cases contained at least 100 mm³ of WMH (Rudie et al., 2019). This finding shows that unlabeled WMH in glioma cohorts is not an isolated occurrence. In joint segmentation settings, these unlabeled WMH regions are implicitly treated as healthy tissue during training. Consequently, the resulting supervision signal may encourage models to misclassify pathological regions as normal anatomy.

When the same voxel space needs to represent healthy brain tissues, tumor lesions, and coexisting abnormalities at the same time, the original BraTS-GLI tumor-subregion labels cannot directly serve as a joint supervision target. This resource gap appears at two levels: coexisting WMH is not systematically represented in the original task, and healthy tissue structures are also not included in the same label space. Table 1 summarizes the data-resource context: established anatomy tools can produce healthy-tissue labels (Fischl, 2012; Billot et al., 2023), and major brain lesion datasets provide single-type tumor or stroke lesion annotations (BraTS 2023 Challenge Organizers, 2023; Absher et al., 2024), but they do not provide healthy-tissue labels, coexisting-lesion labels, and auditable provenance records together in a single reusable data resource. Examples of T1/FLAIR, original BraTS labels, added lesion labels, and unified healthy tissue-lesion labels are shown in Figure 1. To address this gap, we introduce GLI-AL, a WMH-aware anatomy-lesion label resource built upon the full BraTS 2023-GLI training cohort. Specifically, we (1) establish WMH-aware stratification across all cases and distinguish purified and extended subsets, (2) construct a unified label space covering both healthy tissues and lesion structures, and (3) provide traceable label-source metadata and quality-control records for all released cases and files. The resulting resource enables joint supervision of normal anatomy and pathological structures, supports label-noise sensitivity analysis, and facilitates case-source-stratified experimental design. To our knowledge, no publicly available BraTS-GLI-derived resource currently integrates healthy-tissue labels, tumor annotations, WMH-aware stratification, and auditable provenance records within a unified label space.

2. Summary

This resource is intended for research on medical image joint segmentation, brain tumor MRI, label noise, and AI-ready datasets. It provides WMH-aware BraTS-GLI derived labels with transparent label sources and case-level metadata. Researchers can conduct joint training and evaluation for healthy tissues and lesions within the same label space, as well as sensitivity analyses stratified by label source and case subset.

The resource corresponds to 1251 four-modal MRI cases from the BraTS 2023-GLI training set. This paper releases only derived labels and does not redistribute MRI images. The labels use NIfTI format and remain aligned with the upstream BraTS case-folder layout, inheriting the upstream standard-space settings: SRI24 template space, 1 mm³ resolution, 240×240×155 voxel dimensions, and skull-stripped images (Baid et al., 2021). The integer codes of the unified labels are: 0 for background, 1 for cortical gray matter (GM), 2 for basal ganglia (BG), 3 for white matter (WM), 4 for Lesion, 5 for ventricle (Ven), 6 for cerebellum (Cer), and 7 for brainstem (BS). This coding places lesion labels and healthy tissue structures into a single supervision target, so downstream models no longer need to switch between incompatible label definitions. Combined with the upstream BraTS masks, users can retain NCR, ED, and ET, and identify voxels in the unified Lesion class that lie outside the upstream whole-tumor mask as the newly added lesion component. The MedNeXt use case in the validation experiments (Isensee et al., 2021; Roy et al., 2023) uses only T1 and FLAIR; this is an experimental setting, not a modality limitation of the label resource.

The full training-set resource is divided into a 394-case purified subset and an 857-case extended subset; the construction process and label sources are described in Methods. The purified subset includes cases with expert-negative WMH conclusions from the Rudie et al. expert WMH annotation resource (Rudie et al., 2019), together with additional WMH-negative cases identified by model screening. The extended subset merges additional lesion constraints from the remaining cases into the unified Lesion class. In addition to the unified labels, 116 cases also provide repair labels with two foreground classes, representing whole tumor and WMH, respectively. These repair labels




Figure 1: Example images and derived labels. The BraTS label image displays the original BraTS tumor subregions merged as a whole-tumor mask. Added lesion denotes a lesion region that is visible on FLAIR, not covered by the original BraTS tumor-task label, and included in the Lesion class in this paper. The unified healthy tissue–lesion label supplements healthy tissue labels for cortical gray matter, basal ganglia, white matter, ventricle, cerebellum, and brainstem.

copy the label resource of Rudie et al. (Rudie et al., 2019). The repair-label cases contain 105,417 WMH repair voxels in total, corresponding to 0.89% of their accompanying whole-tumor foreground volume, with a median repaired WMH volume of 448.5 voxels per case. These repair labels allow users, after obtaining the upstream BraTS MRI, to reproduce the repaired image inputs used in this study with the accompanying code. Users need to pair this resource’s labels with upstream BraTS 2023-GLI MRI. Apart from Synapse CLI, Python, NiBabel/SimpleITK, or equivalent medical-image I/O tools, no additional registration or resampling is required.

To quantify the amount of information added by this resource, we compared each released unified label with the corresponding original BraTS-GLI expert lesion mask after mapping both files to the same case identifier. The original mask was treated as the pre-existing tumor-task foreground. In contrast, the released unified label contains both healthy-tissue structures and the final Lesion class. This comparison shows that the resource adds foreground labels outside the original expert lesion mask for all 1251 cases. These added labels occupy 1.51 billion voxels, equal to 92.7% of the released foreground. Most of this added foreground comes from healthy-tissue structures that were absent from the original tumor-task label. The lesion-specific addition is smaller but directly addresses the label-noise problem: 857 cases contain additional Lesion voxels outside the original expert lesion foreground, totaling 2.20 million voxels. These voxels represent 1.8% of the final Lesion class and mark coexisting lesion constraints that would otherwise be treated as healthy tissue in a joint segmentation target.

From the perspective of AI-readiness and FAIR use, version v1.0.0 provides controlled-access derived labels and metadata through Synapse project ID syn75210889. Labels are released as NIfTI files in BraTS-compatible case folders, and case identifiers can be aligned with upstream

BraTS 2023-GLI MRI. The release metadata record the stratification composition, case-level provenance and QC status, internal-ID mapping for probability-map generation and label fusion, access conditions, file sizes, and SHA-256 digests; the repository also provides a Synapse CLI download entry point. These metadata allow users to distinguish case sources, label sources, and release boundaries in training, evaluation, and manuscript reporting, rather than treating all derived labels as homogeneous samples. Users must first obtain access to the upstream BraTS project syn51156910, then apply for access to syn75210889 through the email workflow in this resource’s wiki and accept the terms for the derived resource.

3. Discussion

3.1 Strengths

The main strength of this resource is that it handles lesion-label noise caused by unlabeled WMH in BraTS-GLI at the data-source level while keeping the resulting decisions auditable. Label-noise identification, purified/extended stratification, an image-repair reproduction entry point, case-level provenance, QC status, and file-checksum records are fixed parts of the release. This organization allows users to select cases by source, reproduce release boundaries, and report training or evaluation data use without treating all derived labels as homogeneous samples.

A second strength is that the resource provides unified eight-class anatomy-lesion labels for the full training set, allowing lesions and healthy anatomical structures to be jointly modeled within one training target. This makes the resource suitable for joint healthy tissue–lesion supervision and label-noise sensitivity analyses without switching between incompatible task labels.

3.2 Known limitations

The purified subset is not equivalent to 394 cases that have all undergone new manual WMH annotation. It consists of both expert-negative samples and model-screened samples, so users should report its screening source. The additional lesion constraints for the extended subset come from the fusion of automatic tools DeepWMH (Liu et al., 2024) and LST-AI (Wiltgen et al., 2024), which may introduce tool bias, missed detections, and limited sensitivity to small lesions. Healthy-tissue labels come from the automatic tool TumorSynth (Wu et al., 2026) and a probabilistic fusion strategy, and they are not equivalent to manual whole-brain anatomical annotations.

This resource systematically handles WMH comorbid lesions; other coexisting abnormalities, such as microbleeds and lacunar infarcts, have not yet been systematically completed. The external WMH validation set contains only T1 and FLAIR, so the use case in this paper cannot prove the value of T1ce and T2 in all downstream tasks.

3.3 Responsible use

GLI-AL is intended only for scientific research and should not be used for direct clinical diagnosis, treatment decisions, or individual risk assessment. Users should report the data tier and label source used, and should state the limitations of automatically generated labels. Any re-identification attempt is inconsistent with the intended use of this resource. If errors are found in the derived labels, version records, or accompanying code, users should contact the resource contact for this paper: xyx200200@stu.xjtu.edu.cn; issues involving original data or upstream terms should be directed to the official BraTS contact. Future versions will investigate extension of joint healthy tissue and lesion labels to more cases.

4. Resource Availability

4.1 Summary statement

This resource provides BraTS-GLI derived labels and case-level metadata for joint healthy tissue–lesion supervision, addressing the problem that the original tumor-subregion labels cannot directly represent the relationship between coexisting WMH and healthy tissue structures. The resource scope and access route are described in Summary and Resource Availability, the stratified construction workflow is described in Methods, and quality control and example validation results are described in Validation.

4.2 Data and code location

The Synapse repository for this resource is <https://www.synapse.org/Synapse:syn75210889/wiki/>,

and the Synapse project ID is syn75210889. The data resource DOI is <https://doi.org/10.7303/SYN75210889>. The accompanying code repository is <https://github.com/xyx200/brats-gli-anatomy-lesion-code>, which releases the processing scripts authored by this study and the MedNeXt/nnUNet unified-label evaluation adaptation files. This code repository does not redistribute BraTS imaging data, MedNeXt model or training code, model weights, DeepWMH, LST-AI, TumorSynth, NiftySeg, or a complete third-party software environment; third-party tools or model implementations used in this workflow should be obtained from their official sources. Resource versions start from v1.0.0; future label or metadata changes will be documented through versioned records.

This derived resource is a controlled-access labels-only resource. Users must first obtain upstream BraTS 2023 data access in the official BraTS 2023 Synapse project syn51156910 (BraTS 2023 Challenge Organizers, 2023) and accept the applicable terms, then apply for access to GLI-AL through the email template in the syn75210889 wiki and accept the terms for this derived resource.

4.3 Potential use cases

This resource is intended only for scientific research. Expected use cases include glioma MRI joint segmentation research under implicit label noise; joint modeling of healthy brain tissues and unified lesions; analysis of training-set cleaning and the impact of label noise; provenance-based stratified training; and, when shared MRI modalities are available, cross-resource validation from the GLI cohort to external research data. These use cases depend on joint labels for healthy tissues and lesions in the same space and cannot be directly completed with the original BraTS tumor-subregion labels.

4.4 Licensing

This resource is built from controlled-access BraTS 2023-GLI data. Users must first complete the data-access application on the official BraTS 2023 Synapse page and obtain access to the original BraTS-GLI data; this derived resource does not replace or bypass that process. This resource is released under *Creative Commons Attribution-NonCommercial 4.0 International* (CC BY-NC 4.0, <https://creativecommons.org/licenses/by-nc/4.0/>), consistent with the upstream non-commercial use conditions, and remains subject to the applicable BraTS 2023 post-Challenge Terms and Conditions; LICENSE.txt is provided in the repository root.

Publications using this resource must comply with the official BraTS citation requirements and retain the data-source statement: “Data used in this publication were obtained as part of the Brain Tumor Segmentation (BraTS)




Figure 2: Construction workflow of GLI-AL

Challenge project through Synapse ID: syn51156910.” Processing scripts authored by this study in the accompanying code repository are released under the MIT License. Unified-label evaluation files adapted from the MedNeXt/nnUNet evaluation package are released under the Apache License, Version 2.0, with upstream copyright notices and descriptions of this study’s modifications retained. This paper uses MedNeXt as a validation model and does not claim to modify or redistribute its model architecture, training implementation, or model weights. Users must not use this resource for direct clinical application, re-identification attempts, or sharing downloaded label files with unapproved individuals.

4.5 Ethical considerations

This resource consists of controlled-access derived labels generated from BraTS de-identified images and labels. It does not add participant recruitment, scanning, clinical-variable collection, or identity-identifying information, so this derived-label organization work does not require a new ethics approval or exemption process. Informed consent, ethics approval, de-identification, data-release authorization, and the final release decision for the original BraTS data are handled by the upstream BraTS/Synapse data workflow. This resource inherits the upstream access restrictions and use conditions and does not replace or bypass the BraTS 2023-GLI data-access process.

This resource is limited to scientific research that complies with laws, research ethics, and upstream BraTS/Synapse terms. It is not intended for direct clinical diagnosis, treatment decisions, or individual risk assessment. Users must not attempt re-identification and must not share downloaded labels or derived metadata with individuals who have not been approved for access to this resource. If suspected data leakage, violation of access conditions, derived-label errors, or version-record errors are found, users should contact the resource contact for this paper, xxy200200@stu.xjtu.edu.cn; issues involving original BraTS data, upstream consent, approval, or licensing terms should be handled by the official BraTS or Synapse

management workflow.

5. Methods

5.1 Data details

The data source is a BraTS 2023-GLI adult brain glioma training-set cohort with 1251 cases. Each case contains T1, T1ce, T2, and T2-FLAIR MRI provided by the upstream BraTS project; these MRI volumes are not redistributed in this resource. The release is labels-only and contains 1367 NIfTI label files: 1251 unified segmentation labels and 116 image-repair labels. The derived unified labels merge tumor and comorbid lesions into Lesion and add six healthy-tissue classes.

Upstream preprocessing information for the released BraTS images is provided in the official CaPTk BraTS Pre-processing Pipeline documentation: https://cbica.github.io/CaPTk/preprocessing_brats.html. This paper generates derived labels based on the official released results and does not redefine or modify the upstream preprocessing workflow.

The release directory preserves the BraTS-compatible per-case folder layout. Cases requiring upstream MRI repair additionally provide repair labels under `repair_labels/modified/`.

Public metadata files summarize released case/session identifiers, subset and image-repair status, case paths and upstream references, label provenance, QC, license/access terms, size and SHA-256 integrity records, and internal ID mappings used for probability-map generation and label fusion. The released case/session identifier is not claimed to be a verified unique natural-person identifier, and internal IDs are not treated as participant identifiers. Age, sex, scanner, magnetic-field strength, scan protocol, and clinical or molecular variables are controlled by the upstream BraTS workflow; fields not officially released are recorded as unavailable and are not inferred by this resource.

Anonymization, informed consent, original ethics approval, and release decisions for the original data are handled by the upstream BraTS/Synapse workflow; this resource

Table 2: Quantitative label content added relative to the original BraTS-GLI expert lesion masks.

Metric	GM	BG	WM	Lesion outside expert mask	Ven	Cer	BS	All new foreground
Cases with new voxels	1251	1251	1251	857	1251	1251	1251	1251
New voxels (million)	685.57	56.86	518.34	2.20	33.06	176.33	41.28	1513.64
Share of released foreground (%)	41.96	3.48	31.73	0.13	2.02	10.79	2.53	92.65

inherits the upstream de-identification status, access restrictions, and use conditions. Case-level provenance records distinguish upstream BraTS tumor labels, expert-negative WMH conclusions and repair-label WMH components from Rudie et al. (Rudie et al., 2019), and automatic labels or probability maps from DeepWMH, LST-AI, and TumorSynth.

The same paired-label comparison also provides a quantitative view of the released label content (Table 2). The original BraTS-GLI expert lesion foreground is fully retained in the final Lesion class. The added Lesion voxels are concentrated in the extended subset by construction, whereas the six healthy-tissue classes are present in all cases. This result clarifies the resource contribution at two scales. At the lesion scale, the release adds explicit comorbid-lesion constraints in 857 cases. At the whole-label scale, the release changes BraTS-GLI from a lesion-only target into a dense anatomy-lesion supervision target with auditable provenance.

5.2 Methods used for the data creation

The overall construction workflow is shown in Figure 2. The following text presents three steps: purified subset construction, extended subset label completion, and unified anatomy-lesion labels. These three steps serve the same resource objective: to organize the originally separated tumor subregions, coexisting WMH information, and healthy tissue structures in BraTS-GLI into a single label space, while retaining case stratification and quality status so that downstream research can directly compare data quality and joint supervision.

Purified subset construction. This study introduces expert WMH annotations for 285 BraTS cases provided by Rudie et al. (Rudie et al., 2019); these cases come from the BraTS 2018 training set and fall within the BraTS 2023-GLI training set used in this paper. We used the original T1 and FLAIR images and expert lesion annotations of these 285 cases to train MedNeXt (Isensee et al., 2021; Roy et al., 2023) in a five-fold configuration, enabling the model to segment tumor regions and coexisting WMH separately. T1 and FLAIR were selected because the target abnormality to

be screened out in this study is WMH: FLAIR presents white matter hyperintensities, and T1 provides complementary information.

For the 233 cases with WMH in the expert annotations, NiftySeg (Cardoso et al., 2012) was used to perform in-place filling of the WMH regions, replacing them with surrounding normal-tissue signal. The trained MedNeXt then screened all candidate cases without an expert-negative conclusion and considered only predicted WMH outside the existing tumor masks. Screening identified 342 model-negative cases. Of these, 116 came from positive cases with expert WMH annotations and had already been repaired; the other 226 came from cases without expert WMH annotations. These 342 cases, together with the 52 expert-negative cases, formed the 394-case purified subset. For the 116 repaired cases, the repaired images replaced the corresponding original images in the cohort and were not counted as duplicate cases. The remaining 857 cases forming the extended subset retained their original imaging status.

Extended subset label completion. For the non-purified 857-case extended subset, the original four-modal images remain unchanged. This paper uses DeepWMH (Liu et al., 2024) and LST-AI (Wiltgen et al., 2024) to generate candidate masks for coexisting abnormalities, and takes their intersection to reduce false positives from any single tool. This intersection mask is united with the existing BraTS lesion mask to form the unified Lesion class.

Unified anatomy-lesion labels. To obtain healthy-tissue labels, TumorSynth (Wu et al., 2026) was run separately on the four MRI modalities to obtain voxel-level softmax probability maps for each modality. The original outputs were aggregated into seven foreground classes through a mapping table: GM, BG, WM, Lesion, Ven, Cer, and BS. We introduced automatic outlier detection based on the interquartile range (IQR): at the whole-dataset level, the total number of foreground voxels predicted by each modality was counted, and predictions below the lower outlier bound $Q_1 - 1.5 \times \text{IQR}$ were marked as outliers. After manual review, low-quality modality probability maps that caused missing foreground were removed to prevent noisy predictions from entering subsequent fusion.

Table 3: Comparison of joint segmentation performance

Dataset	Method	Metric	GM	BG	WM	Lesion	Ven	Cer	BS
GLI (in-domain)	Baseline	DSC↑	96.0 (1.9)	95.2 (3.9)	95.9 (2.2)	90.5 (9.8)	94.8 (3.8)	98.4 (2.9)	97.2 (3.4)
		HD95↓	1.07 (0.29)	1.21 (0.64)	1.10 (0.36)	15.71 (18.47)	1.25 (0.83)	1.21 (1.89)	1.22 (1.70)
	Joint-Baseline	DSC↑	96.3 (1.6)	95.9 (3.1)	96.2 (1.9)	90.8 (10.3)	95.7 (2.8)	98.5 (2.8)	97.6 (3.1)
		HD95↓	1.05 (0.23)	1.11 (0.43)	1.09 (0.30)	18.90 (30.06)	1.11 (0.41)	1.19 (1.88)	1.21 (1.70)
WMH (out-of-domain)	Baseline	DSC↑	79.1 (2.1)	88.3 (2.0)	89.8 (2.7)	31.1 (27.5)	89.1 (3.5)	91.4 (1.5)	87.4 (2.1)
		HD95↓	6.55 (1.34)	1.67 (0.32)	1.65 (0.35)	140.40 (150.46)	1.84 (0.48)	4.84 (0.99)	2.31 (0.49)
	Joint-Baseline	DSC↑	79.3 (2.1)	88.4 (2.0)	89.4 (3.2)	4.3 (12.9)	89.6 (3.4)	91.3 (1.6)	87.1 (2.2)
		HD95↓	6.50 (1.33)	1.68 (0.32)	1.92 (0.64)	297.46 (132.37)	1.77 (0.48)	4.84 (1.01)	2.34 (0.51)

Let the lesion hard constraint at voxel $\mathbf{v}$ be $M(\mathbf{v})$, and let the lesion class index be c_l . For the probability map of valid modality m , the correction is:

$$P'_m(c | \mathbf{v}) = \begin{cases} 1, & M(\mathbf{v}) > 0 \text{ and } c = c_l, \\ 0, & M(\mathbf{v}) > 0 \text{ and } c \neq c_l, \\ 0, & M(\mathbf{v}) = 0 \text{ and } c = c_l, \\ \frac{P_m(c | \mathbf{v})}{\sum_{k \neq c_l} P_m(k | \mathbf{v})}, & M(\mathbf{v}) = 0 \text{ and } c \neq c_l. \end{cases} \quad (1)$$

After the lesion channel is removed outside the lesion region, the corrected probabilities are renormalized. The voxel-wise uncertainty and adaptive weight of each valid modality are:

$$H_m(\mathbf{v}) = - \sum_c P'_m(c | \mathbf{v}) \log P'_m(c | \mathbf{v}), \quad (2)$$

$$W_m(\mathbf{v}) = \frac{\exp[-H_m(\mathbf{v})]}{\sum_{j=1}^N \exp[-H_j(\mathbf{v})]}, \quad (3)$$

where N denotes the number of valid modalities retained after quality control, and entropy is computed using the conventional definition $0 \log 0 = 0$. The final fused probability and integer label are:

$$L_{\text{fused}}(\mathbf{v}) = \arg \max_c \sum_{m=1}^N W_m(\mathbf{v}) P'_m(c | \mathbf{v}). \quad (4)$$

The labels generated by Eqs. (1) to (4) fix the unified lesion constraint and preserve high-confidence predictions of healthy anatomical structures from each modality.

6. Validation

6.1 Description of approaches ensuring quality of data

Quality control corresponds to four levels of the data-creation workflow and release integrity. First, expert-negative samples in the purified subset come from cases with strictly zero WMH volume among the 285 expert WMH annotations; the remaining purified cases were selected as high-confidence WMH-negative samples by five-fold MedNeXt trained from the expert annotations. Second, the additional lesion constraints for the extended

subset use the intersection of DeepWMH and LST-AI and are united with existing BraTS lesion masks, reducing false positives from a single tool while retaining the original target lesions. Third, the four-modal TumorSynth probability maps undergo IQR outlier detection based on total foreground voxel counts, and outlier modalities are removed after manual review before entropy-weighted fusion. Fourth, after freezing the release tree, labels-only case-level and file-level integrity checks were completed: case counts and subset counts, unified-label and repair-label paths, file roles, file sizes, SHA-256 digests, access conditions, and manifest alignment were checked. A total of 1251/1251 cases passed and 0 failed; all 1367 NIfTI label files entered the release manifest, including 1251 unified segmentation labels and 116 image-repair labels, and the number of redistributed MRI images was 0.

6.2 Controlled ML validation

This paper uses a controlled ML use case to answer a resource-level question: in a unified anatomy-lesion label space, why is it necessary to construct and release a purified subset separately, rather than relying only on more noisy training cases? All experiments use MedNeXt-M (Isensee et al., 2021; Roy et al., 2023) with default settings: $3 \times 3 \times 3$ kernel, five-fold cross-validation, AdamW optimizer, initial learning rate 1×10^{-3} , weight decay 3×10^{-5} , 1000 epochs, and batch size 2. Baseline was trained only on the 394-case purified subset. The 1000 training cases for Joint-Baseline consisted of the 394-case purified subset and 606 additional cases randomly selected from the extended subset, and the latter did not receive additional comorbid-lesion labels. Both experiments used T1 and FLAIR inputs.

For each foreground class, this paper uses the Dice similarity coefficient (DSC) to measure the overlap between the predicted mask and the reference mask. When both the predicted mask and the reference mask are empty, DSC is recorded as 1.0. HD95 is the 95th percentile of the surface distance between the predicted boundary and the reference boundary; smaller values indicate closer boundaries. It is recorded as 0.0 when both masks are empty. If only one side




Figure 3: Visualization analysis of in-domain and out-of-domain joint segmentation

is empty, if the abnormal prediction fills the whole image, or if distance computation fails or exceeds the threshold, the maximum distance penalty value 374.0 is assigned. Metrics in the table are first computed within each case and then reported as case-level mean (standard deviation).

GLI in-domain testing used the remaining 251 cases. External WMH testing used 170 images from the MICCAI 2017 WMH dataset, which originally provides T1, FLAIR, and lesion masks (Kuijf et al., 2019). Following the BraTS processing workflow, this paper used HD-BET for skull stripping (Isensee et al., 2019), used ANTs (Tustison et al., 2021) to rigidly register T1 to SRI24 (Rohlfing et al., 2010), and applied the same transform to FLAIR and lesion labels. Healthy-tissue labels were generated by WMH-SynthSeg (Laso et al., 2024) and mapped to the same seven foreground classes as GLI. The same lesion hard constraint and entropy-weighted fusion rules as in the GLI data were then used to generate the joint reference labels. This external validation set is not part of the GLI resource released in this paper.

Table 3 shows that, in GLI in-domain evaluation, the healthy-tissue DSC and HD95 of Baseline and Joint-Baseline are close, indicating that a smaller amount of high-quality data can still maintain stable healthy-brain-tissue segmentation ability. The lesion DSC of Joint-Baseline is only slightly higher than that of Baseline, but its HD95 is larger and has a higher standard deviation, indicating that the two models have similar overlap performance for the main lesion, while Joint-Baseline is more likely to miss WMH that is far from the main lesion and smaller in volume, thereby increasing distance error. In external zero-shot WMH evaluation, the two models still maintain similar results for healthy-tissue classes, indicating that model performance on out-of-domain healthy anatomical-structure segmentation is

preserved, whereas Baseline has clearly higher lesion DSC and lower HD95 than Joint-Baseline. This result explains the resource value of the purified subset: it provides a WMH-aware cleaned training reference for the joint label space, allowing researchers to distinguish the effects of data quality and sample size on sensitivity to coexisting lesions. The relatively low overall out-of-domain GM metric mainly comes from differences in reference-label generation protocols: healthy tissues in the GLI training labels are generated by TumorSynth, whereas WMH test labels are generated by WMH-SynthSeg. Their gray-matter boundary definitions are not fully consistent, so this phenomenon should not be interpreted as a failure of healthy-tissue segmentation ability. The qualitative results in Figure 3 are consistent with the table: in the GLI sample, both models cover the main tumor region, and the differences mainly appear in the segmentation of implicit WMH; in the WMH sample, Baseline detects more WMH regions aligned with GT, whereas Joint-Baseline shows no response to WMH lesions.

7. Conflicts of Interest

The authors declare that they have no conflicts of interest.

8. Acknowledgements

The authors thank the organizers of the Brain Tumor Segmentation (BraTS) Challenge and the data-providing institutions for establishing and releasing de-identified brain MRI research resources. Data used in this publication were obtained as part of the Brain Tumor Segmentation (BraTS) Challenge project through Synapse ID: syn51156910. The authors also thank Rudie et al. for the publicly available

expert annotation materials that supported the construction of this derived label resource. This work was supported in part by the National Natural Science Foundation of China (Nos. T2522028 and 12326616), Natural Science Basic Research Program of Shaanxi (No. 2024JC-TBZC-09), and Shaanxi Provincial Key Industrial Innovation Chain Project (No. 2024SF-ZDCYL-02-10).

References

- John Absher, Sarah Goncher, Roger Newman-Norlund, Nicholas Perkins, Grigori Yourganov, Jan Vargas, Sanjeev Sivakumar, Naveen Parti, Shannon Sternberg, Alex Teghipco, Makayla Gibson, Sarah Wilson, Leonardo Bonilha, and Chris Rorden. The Stroke Outcome Optimization Project: Acute ischemic strokes from a comprehensive stroke center. *Scientific Data*, 11(1):839, 2024. doi: 10.1038/s41597-024-03667-5.
- Ujjwal Baid, Satyam Ghodasara, Suyash Mohan, Michel Bilello, Evan Calabrese, Erwin Colak, Keyvan Farahani, Jayashree Kalpathy-Cramer, Justin D. Kirby, Federico R. Kitamura, et al. The rsna-asnr-miccai brats 2021 benchmark on brain tumor segmentation and radiogenomic classification. *arXiv preprint arXiv:2107.02314*, 2021. doi: 10.48550/arXiv.2107.02314. URL <https://arxiv.org/abs/2107.02314>.
- Spyridon Bakas, Hamed Akbari, Aristeidis Sotiras, Michel Bilello, Martin Rozycki, Justin S. Kirby, John B. Freymann, Keyvan Farahani, and Christos Davatzikos. Advancing the cancer genome atlas glioma mri collections with expert segmentation labels and radiomic features. *Scientific Data*, 4(1):170117, 2017a. doi: 10.1038/sdata.2017.117.
- Spyridon Bakas, Hamed Akbari, Aristeidis Sotiras, Michel Bilello, Martin Rozycki, Justin S. Kirby, John B. Freymann, Keyvan Farahani, and Christos Davatzikos. Segmentation labels and radiomic features for the pre-operative scans of the tcga-gbm collection, 2017b. URL <https://doi.org/10.7937/K9/TCIA.2017.KLXWJJ1Q>.
- Spyridon Bakas, Hamed Akbari, Aristeidis Sotiras, Michel Bilello, Martin Rozycki, Justin S. Kirby, John B. Freymann, Keyvan Farahani, and Christos Davatzikos. Segmentation labels and radiomic features for the pre-operative scans of the tcga-lgg collection, 2017c. URL <https://doi.org/10.7937/K9/TCIA.2017.GJQ7R0EF>.
- Benjamin Billot et al. SynthSeg: Segmentation of brain MRI scans of any contrast and resolution without retraining. *Medical Image Analysis*, 86:102789, 2023. doi: 10.1016/j.media.2023.102789.
- BraTS 2023 Challenge Organizers. BraTS 2023 Challenge Synapse Project Page. <https://www.synapse.org/Synapse:syn51156910/wiki/621282>, 2023. Synapse project syn51156910, wiki page 621282; last modified 2025-08-28; accessed 2026-06-19.
- M. Jorge Cardoso, Matthew J. Clarkson, Marc Modat, and Sébastien Ourselin. NiftySeg: open-source software for medical image segmentation, label fusion and cortical thickness estimation. *IEEE International Symposium on Biomedical Imaging*, 2012.
- Bruce Fischl. FreeSurfer. *NeuroImage*, 62(2):774–781, 2012. doi: 10.1016/j.neuroimage.2012.01.021.
- Fabian Isensee, Marianne Schell, Irada Pflueger, Gianluca Brugnara, David Bonekamp, Ulf Neuberger, Antje Wick, Heinz-Peter Schlemmer, Sabine Heiland, Wolfgang Wick, et al. Automated brain extraction of multisequence MRI using artificial neural networks. *Human Brain Mapping*, 40(17):4952–4964, 2019. doi: 10.1002/hbm.24750.
- 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.
- Alexandros Karargyris, Renato Umeton, Micah J. Sheller, Alejandro Aristizabal, Jithin George, Andreas Wuest, Sarthak Pati, Hossam Kassem, Maximilian Zenk, Daniel J. Beutel, et al. Federated benchmarking of medical artificial intelligence with medperf. *Nature Machine Intelligence*, 5: 799–810, 2023. doi: 10.1038/s42256-023-00652-2. URL <https://doi.org/10.1038/s42256-023-00652-2>.
- Hugo J. Kuijf, J. Matthijs Biesbroek, Jeroen De Bresser, Rutger Heinen, Simon Andermatt, Mariana Bento, Matt Berseth, Mikhail Belyaev, M. Jorge Cardoso, Adrià Casamitjana, et al. Standardized assessment of automatic segmentation of white matter hyperintensities and results of the WMH segmentation challenge. *IEEE Transactions on Medical Imaging*, 38(11):2556–2568, 2019. doi: 10.1109/TMI.2019.2905770.
- Pedro Laso, Stefano Cerri, Amy Sorby-Adams, et al. Quantifying white matter hyperintensity and brain volumes in heterogeneous clinical and low-field portable MRI. In *2024 IEEE International Symposium on Biomedical Imaging (ISBI)*, pages 1–5. IEEE, 2024. doi: 10.1109/ISBI56570.2024.10635502.

Chenghao Liu, Zhiyong Zhuo, Liying Qu, Ying Jin, Tiantian Hua, Jun Xu, Guirong Tan, Yuna Li, Yunyun Duan, Tingting Wang, et al. DeepWMH: A deep learning tool for accurate white matter hyperintensity segmentation without requiring manual annotations for training. *Science Bulletin*, 69(7):872–875, 2024. doi: 10.1016/j.scib.2024.01.034.

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, 2015. doi: 10.1109/TMI.2014.2377694.

Torsten Rohlfing, Natalie M. Zahr, Edith V. Sullivan, and Adolf Pfefferbaum. The sri24 multichannel atlas of normal adult human brain structure. *Human Brain Mapping*, 31(5):798–819, 2010. doi: 10.1002/hbm.20906.

Saikat Roy, Gregor Koehler, Constantin Ulrich, Michael Baumgartner, Jens Petersen, Fabian Isensee, Paul F. Jaeger, and Klaus H. Maier-Hein. MedNeXt: Transformer-driven scaling of convnets for medical image segmentation. In *International Conference on Medical Image Computing and Computer-Assisted Intervention (MICCAI)*, pages 405–415. Springer, 2023. doi: 10.1007/978-3-031-43901-8_39.

Jeffrey D. Rudie, David A. Weiss, Rachit Saluja, Andreas M. Rauschecker, Jiancong Wang, Leo Sugrue, Spyridon Bakas, and John B. Colby. Multi-disease segmentation of gliomas and white matter hyperintensities in the brats data using a 3d convolutional neural network. *Frontiers in Computational Neuroscience*, 13:84, 2019. doi: 10.3389/fncom.2019.00084.

Nicholas J. Tustison, Philip A. Cook, Andrew J. Holbrook, Hans J. Johnson, John Muschelli, Gabriel A. Devenyi, Jeffrey T. Duda, Sandhitsu R. Das, Nicholas C. Cullen, Daniel L. Gillen, et al. The ANTsX ecosystem for quantitative biological and medical imaging. *Scientific Reports*, 11(1):9068, 2021. doi: 10.1038/s41598-021-87564-6.

T. Wiltgen, J. McGinnis, S. Schlaeger, et al. LST-AI: A deep learning ensemble for accurate MS lesion segmentation. *NeuroImage: Clinical*, 42:103611, 2024. doi: 10.1016/j.nicl.2024.103611.

Jian Wu, Benjamin Billot, Fang Zhao, et al. TumorSynth: Integrated brain tumor and tissue segmentation on brain mri scans of any resolution and contrast. *Radiology: Imaging Cancer*, 8(2):e250222, 2026. doi: 10.1148/rycan.250222.