This MICCAI paper is the Open Access version, provided by the MICCAI Society. It is identical to the accepted version, except for the format and this watermark; the final published version is available on SpringerLink.

Anatomy-Aware Healthy Brain Inpainting: Integrating Tissue Guidance Across Inputs, Features, and Objectives

Hanxue Gu, Yan Bao, Kaizar Rangwala, Kang Wang, and Yang Yang

University of California, San Francisco, CA, USA
hanxue.gu@ucsf.edu

Abstract. Healthy brain inpainting replaces tumor-affected regions in three-dimensional MRI with healthy tissue while exactly preserving observed voxels outside the missing region. Because the missing anatomy is not uniquely determined by its surroundings, a model must balance similarity to the hidden reference with plausible tissue organization. We present our BraTS 2026 Task 4 system and study anatomical guidance in three forms: an additional anatomical input, an auxiliary tissue-prediction task, and a segmentation-based training loss. All variants use a mask-conditioned 3D U-Net that copies the observed image outside the fill region. The results show that the value of anatomical guidance depends on how it is introduced. Providing a completed anatomical label map as an input did not improve image fidelity over mask-only conditioning. Jointly predicting major tissue classes improved the structural-similarity index (SSIM) by 0.0082 for the largest missing regions. Adding this segmentation-based loss increased agreement with automatically derived anatomical labels by 1.58% overall and by 3.95% in voids of at least 45,000 voxels, while image-fidelity metrics remained broadly stable. The submitted direct model achieved 0.8109 SSIM, 0.00737 mean squared error, and 22.65 dB peak signal-to-noise ratio on the official validation set. These findings show that anatomical objectives are most useful for improving tissue organization in large missing regions, but they should be evaluated alongside reference-based fidelity because the two criteria capture different properties of a reconstruction.

Keywords: Brain MRI · Image inpainting · Healthy tissue synthesis · Anatomical guidance · Segmentation consistency · BraTS

1 Introduction

Healthy brain inpainting reconstructs tissue within a prescribed region of a voided magnetic resonance image while leaving the observed brain unchanged. The task can support registration, atlas-based analysis, parcellation, and other pipelines that assume approximately healthy anatomy. In the BraTS Local Synthesis setting, the input is a T1-weighted volume with a masked healthy region and a tumor region, and the desired output is a complete volume in which both

regions are replaced by healthy tissue [3,5]. Unlike general image completion, this setting combines a counterfactual target inside the lesion with a strict identity requirement outside the fill mask.

The central difficulty is that the hidden anatomy is only partially determined by the visible context. A small void may be constrained by nearby cortical and ventricular boundaries, whereas a large void can admit several anatomically reasonable completions. The challenge nevertheless compares the prediction with one hidden subject-specific reference using structural similarity (SSIM), mean squared error (MSE), and peak signal-to-noise ratio (PSNR). Direct regression is well suited to this fidelity objective, and deterministic 3D U-Nets form a strong family for the task [11,12]. However, an intensity estimate that matches the reference well is not guaranteed to exhibit readily parseable gray-matter, white-matter, and cerebrospinal-fluid structure.

We study anatomical guidance as a design space rather than as a single added module. At the *input level*, completed parcellations supply an explicit anatomical hypothesis through label modulation or derived image channels. At the *feature level*, a four-class tissue head learns background, cerebrospinal fluid (CSF), gray matter (GM), and white matter (WM), with optional soft feedback to the image decoder. At the *objective level*, a frozen 32-class segmenter supplies differentiable consistency gradients without changing the inference inputs. Only input-level conditioning requires a completed parcellation at inference; the feature- and objective-level designs retain the direct image-and-mask interface.

The experiments reveal complementary behavior across these designs. Completed parcellations do not outperform mask-only conditioning, while feature- and objective-level guidance requires no additional anatomical input at inference. Coarse-tissue auxiliary learning yields its largest gain in the largest-void quartile. Frozen-segmenter consistency reduces absolute intensity error on large voids, yet lowers SSIM as the void grows. Together, these results show that anatomy-aware objectives change which plausible completion is selected and that this change cannot be summarized by one image metric.

Our contributions are:

1. an identity-preserving 3D inpainting system evaluated under the official challenge contract;
2. a unified study of input-, feature-, and objective-level anatomical guidance; and
3. a void-size analysis of the SSIM–MSE/PSNR trade-off, with complementary surrogate-label and coarse-tissue Dice.

2 Method

2.1 Overview and Problem Formulation

Let $x_v \in \mathbb{R}^{H\times W\times D}$ denote a voided T1-weighted image, $m \in \{0,1\}^{H\times W\times D}$ the complete fill mask, and y the healthy training target. The mask is the union of

the healthy void and the tumor region. The inpainting network predicts an image proposal $\tilde{y} = f_\theta([x_v, m])$, and the submitted image is

$$\hat{y} = m \odot \tilde{y} + (1 - m) \odot x_v. \tag{1}$$

Equation (1) prevents the model from modifying observed voxels. We denote the healthy subset by Ω_h. Only this region has a valid healthy target, because the observed tumor depicts pathology rather than the desired counterfactual tissue.

Figure 1 summarizes the three locations at which anatomy is introduced. Every model uses hard output composition, and the controlled feature- and objective-level arms use the same direct inpainting backbone. Fine anatomical labels contain 32 SynthSeg-derived classes [1]; coarse tissue labels contain four classes (background, CSF, GM, and WM). Keeping these representations separate is important because they serve different purposes.

2.2 Anatomical Labels and Training Voids

We apply SynthSeg v1 directly to each complete, unnormalized T1-weighted training volume and verify that the output shape and affine match the source image [1]. Its 32 retained labels are remapped to contiguous indices. For the coarse target, ventricular and CSF labels {4,5,14,15,43,44} form CSF; white-matter labels {2,7,16,41,46} form WM; and the remaining non-background labels form GM. Only the known healthy void supervises anatomy, because tumor-region labels do not describe the desired counterfactual tissue.

Each case supplies five training voids. One is the provided healthy mask; four are transformed healthy-mask shapes sampled from the cohort-wide mask pool and placed inside the brain, separated from the tumor, and mutually non-overlapping. The selected healthy void is unioned with the case's tumor mask to define the complete fill mask, which also determines the crop center.

2.3 Direct Inpainting Backbone

The direct backbone is a three-level 3D U-Net reproduced from the released BraTS-winning family [11,12]. It uses repeated 3×3×3 convolution–normalization–activation blocks, stride-2 encoding, skip-connected decoding, and a 1×1×1 output layer with tanh activation. Inputs use a fixed $208 \times 208 \times 144$ crop centered on the fill mask. For PM1 preprocessing, negative intensities are clipped, each voided crop is divided by its maximum, and the result is mapped to $[-1, 1]$. The baseline objective is

$$\mathcal{L}_{\text{base}} = \text{MAE}_{\Omega_h}(\hat{y}, y) + \mathcal{L}_{\text{SSIM}}(\hat{y}, y), \tag{2}$$

where the MAE term uses only the healthy void and the SSIM loss is evaluated over the crop. The two terms have equal weight. SSIM supplies a local luminance, contrast, and structure objective [9], while the masked MAE directly constrains valid target voxels.

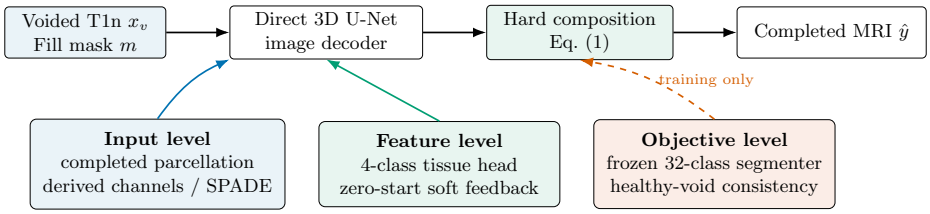

Fig. 1. Conceptual overview of anatomy integration. A shared direct U-Net predicts intensities that are composited with the observed image to guarantee identity outside the fill mask. Anatomy is investigated as an explicit fine-label prior, a coarse-tissue feature branch, or a frozen-segmenter objective used only during training. Colors encode integration level rather than measured performance.

2.4 Input-Level Anatomy and Mask Conditioning

The input-level screen compares two uses of a completed 32-class parcellation. SPADE transforms its one-hot representation into spatially varying affine parameters within decoder normalization layers [7]. The concatenation arm converts the same labels into a subject-calibrated reference image: each class receives its mean intensity from observed tissue, and the completed label map assigns those values inside the void. This reference and the contralateral T1 image are appended to the voided image and fill mask. For both arms, observed SynthSeg labels are retained outside the void; inside, the contralateral label is used when present and a population atlas otherwise.

We separately evaluate mask-gated convolution as a spatial-conditioning control [10]. The binary fill mask modulates learned features throughout the network, but it supplies no tissue labels and is therefore not an anatomy-guided arm.

2.5 Coarse-Tissue Auxiliary Feedback

Coarse-tissue feedback learns anatomical organization without requiring a hard 32-class input. Let z be the final decoder representation. An auxiliary head predicts four tissue logits $q(z)$, while the image head predicts $a(z)$. Optional feedback is

$$\tilde{y} = \tanh\big[a(z) + \alpha\, b\big(\mathrm{softmax}(q(z)) - \tfrac{1}{4}\big)\big]\,, \tag{3}$$

where b is a bias-free $1 \times 1 \times 1$ adapter initialized to zero. The centering and zero initialization preserve the pretrained image mapping at the start of training.

Three arms isolate the contribution of tissue supervision and feedback: **A** uses image reconstruction alone; **B** adds the tissue loss but sets $\alpha = 0$; and **C** adds the same tissue loss with $\alpha = 1$. The healthy-void tissue objective is

$$\mathcal{L}_{\mathrm{tissue}} = \mathcal{L}_{\mathrm{CE}} + 0.5\,\mathcal{L}_{\mathrm{macroDice}}. \tag{4}$$

Samples are weighted by healthy-void quartile, emphasizing cases in which tissue organization must be inferred over a wider region. The quartile weights are $(0.5, 0.75, 1.0, 1.75)$. The feature-level objective is $\mathcal{L}_{\mathrm{base}} + 0.03 w_q \mathcal{L}_{\mathrm{tissue}}$.

2.6 Frozen-Segmenter Consistency

Objective-level guidance provides fine-anatomy gradients without modifying the inference inputs. We use a frozen 32-class 3D U-Net S. The complete objective is

$$\mathcal{L} = \mathcal{L}_{\text{base}} + \lambda\, \mathbf{1}\{|\Omega_h| \geq 45{,}000\}\ \mathrm{CE}(S(\hat{y})_{\Omega_h}, \ell_{\Omega_h})\,, \tag{5}$$

where ℓ is the SynthSeg-derived fine parcellation and $\lambda \in \{0, 0.01, 0.03\}$. The segmenter weights are fixed, but gradients flow through S into the inpainting network. Supervision is limited to the healthy void because labels under the tumor do not describe the desired counterfactual tissue.

The three weights span no, weak, and stronger anatomical regularization. The 45,000-voxel threshold focuses the additional objective on large voids and defines the corresponding evaluation stratum.

Because the two networks use different intensity conventions, a differentiable robust-percentile bridge converts the inpainting output to the segmenter's input range while preserving gradient flow.

2.7 Optimization and Inference

Models use batch size one, Adam at 10^{-4}, and exponential moving-average weights with decay 0.9995. Training uses a patient-grouped five-fold split and the voids described in Section 2.2.

The objective-weight ablation varies only λ. At inference, the EMA model performs one direct forward pass on the centered crop, applies Eq. (1), and restores the original NIfTI geometry. Sliding-window tiling is not used.

3 Experiments and Results

3.1 Data and Evaluation Protocol

The BraTS Local Synthesis training set contains 1,251 cases [3]. Models are evaluated on a patient-grouped held-out fold; the input-conditioning screen uses a fixed development subset of this fold. Official validation performance is scored separately on Synapse.

All quantitative evaluation uses version 0.0.12 of the official `inpainting` package [4]. The evaluator receives the prediction, target, healthy-void mask, and original voided volume as the normalization reference. The primary metrics are SSIM (higher is better), MSE (lower is better), and PSNR (higher is better). For the segmentation-consistency study we use paired per-case differences and a two-sided Wilcoxon signed-rank test. Healthy voids are stratified as small ($<$ 10k voxels), medium (10k–45k), large ($\geq$ 45k), with the largest group further split at 90k voxels. We also apply the frozen segmenter to each completion and compare its hard labels with SynthSeg-derived labels inside the healthy void. $\mathrm{SegDice}_{\Omega_h}$ is foreground macro Dice over reference-present classes in this fixed surrogate label space; it is not agreement with manual neuroanatomical annotations.

3.2 Anatomy Guidance Across Integration Levels

Table 1 summarizes how each integration level changes image fidelity or tissue consistency relative to its corresponding baseline.

Table 1. Summary of anatomy-guidance effects. Each row reports the change produced by adding the indicated anatomy mechanism to its corresponding baseline; negative MSE is favorable.

Level	Comparison	ΔSSIM	ΔMSE ($\times 10^{-3}$)	Complementary change
Input	Completed-label SPADE − mask-gated	−0.002	+1.470	−0.589 dB PSNR
Feature	Tissue + feedback − image-only	+0.001	+0.150	+0.008 Q4 SSIM
Objective	Seg. consistency − control	+0.001	+0.140	+1.580% SegDice

We first examine whether an explicit parcellation hypothesis improves image synthesis. Table 2 reports the completed-parcellation screen. SPADE conditioning uses observed SynthSeg labels, contralateral context, and an atlas. Relative to mask-gated conditioning, SSIM changes by −0.002, MSE by +0.001, and PSNR by −0.589 dB. The completed anatomical hypothesis therefore does not improve image fidelity in this input-level comparison.

Table 2. Input-level anatomy development screen. Explicit parcellation conditioning does not improve the official image-fidelity metrics.

Input condition	SSIM ↑	MSE ($\times 10^{-3}$) ↓	PSNR ↑
Mask-gated image + fill mask	**0.658**	**11.890**	**19.590**
Completed fine labels via SPADE	0.656	13.350	19.000

At the feature level, coarse-tissue learning avoids this hard commitment. Table 3 reports the controlled A/B/C study. Both tissue-supervised arms improve the largest-quartile SSIM by approximately 0.008 relative to the image-only arm. Both tissue-supervised arms reach 0.820 aggregate SSIM, and their small difference indicates that most of the improvement is associated with learning the auxiliary tissue task.

Table 3. Coarse-tissue auxiliary-feedback ablation. Q4 denotes the largest healthy-void quartile and is reported relative to A.

Arm	Configuration	SSIM ↑	MSE ($\times 10^{-3}$) ↓	PSNR_{01} ↑	ΔQ4 SSIM ↑
A	Image-only	0.819	**7.810**	**22.746**	—
B	Tissue auxiliary	**0.820**	8.100	22.439	**+0.008**
C	Tissue + feedback	**0.820**	7.960	22.513	**+0.008**

3.3 Segmenter Consistency Reallocates Error

Table 4 summarizes the final objective-level comparison. Segmentation consistency raises 32-class SegDice by 1.580% across all cases and by 3.954% in the gated large-void stratum. Coarse-tissue Dice likewise rises from 75.20% to 76.82% on gated cases, while cohort-level image-fidelity metrics remain close between methods. Figure 2 examines the corresponding case-level trade-off.

Table 4. Effect of the segmentation-consistency objective. Δ is segmentation consistency minus control. Negative MSE and positive SSIM/PSNR/SegDice are favorable. p is a paired two-sided Wilcoxon test.

Metric	Stratum	Control	Seg.	Δ	Win (%)	p
SSIM	All	0.821	0.822	+0.001	53.2	5.2×10^{-2}
	$<$ 10k	0.942	0.943	+0.001	51.1	4.3×10^{-1}
	10–45k	0.777	0.780	+0.003	57.9	1.8×10^{-2}
	$\geq$ 45k	0.725	0.725	0.000	51.2	8.9×10^{-1}
	$\geq$ 90k	0.690	0.689	−0.001	50.0	9.6×10^{-1}
MSE ($\times10^{-3}$)	All	8.020	8.160	+0.140	57.2	1.6×10^{-3}
	$\geq$ 45k	12.360	12.640	+0.280	67.1	4.3×10^{-3}
PSNR_{01}	All	22.974	23.051	+0.077	57.2	2.0×10^{-3}
	$\geq$ 45k	20.371	20.474	+0.103	67.1	2.3×10^{-3}
$\text{SegDice}_{\Omega_h}$ (%)	All	64.222	65.802	+1.580%	65.2	1.5×10^{-10}
	$\geq$ 45k	59.812	63.767	+3.954%	84.1	4.0×10^{-10}

Table 5 condenses the final paired comparison used in Figure 2 and retains the objective-weight sensitivity. Segmentation consistency improves aggregate SSIM, PSNR_{01}, and SegDice in the final pair. Under the shared sensitivity budget, $\lambda = 0.03$ achieves the highest SegDice and PSNR_{01}.

Table 5. Final paired comparison and objective-weight sensitivity. The top block is the final pair analyzed in Figure 2; the bottom block compares objective weights under a shared training budget. Δ is segmentation consistency minus control. SegDice is measured in the fixed SynthSeg label space. Bold indicates the best value within each block.

Model	SSIM ↑	MSE ($\times10^{-3}$) ↓	PSNR_{01} ↑	$\text{SegDice}_{\Omega_h}$ (%) ↑
Final paired comparison				
Control, $\lambda = 0$	0.821	**8.020**	22.974	64.222
Seg. consistency, $\lambda = 0.01$	**0.822**	8.160	**23.051**	**65.802**
Δ	+0.001	+0.140	+0.077	+1.580%
Objective-weight sensitivity				
Control, $\lambda = 0$	**0.819**	**7.380**	22.990	64.650
Seg. consistency, $\lambda = 0.01$	**0.819**	7.940	22.970	65.220
Seg. consistency, $\lambda = 0.03$	**0.819**	7.770	**23.050**	**66.730**

Figure 2 exposes the void-size dependence directly. For the final $\lambda = 0.01$ model, mean SSIM changes by +0.00062, +0.00276, +0.00037, and −0.00096 across successively larger void strata. Thus, the mean SSIM benefit is concentrated below 90k voxels. SegDice improves in 65.2% of cases; joint SegDice–SSIM improvement occurs in 38.8%, whereas SegDice improves with lower SSIM in 26.4%. This variation is hidden by the aggregate mean.

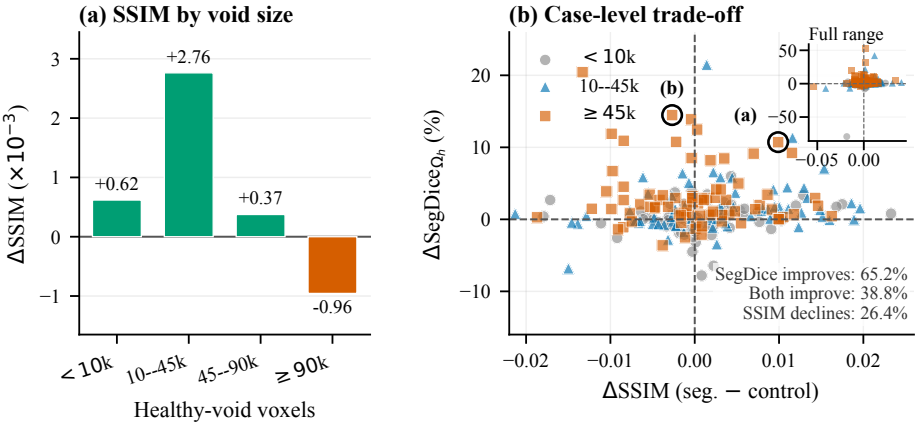

Fig. 2. Anatomical consistency and image fidelity have case-dependent effects. **(a)** Mean SSIM difference by void size for the final $\lambda = 0.01$ model. **(b)** Paired changes for the same comparison; SegDice changes are shown in percent, colors encode healthy-void size, and the inset shows the full range. Outlined points correspond to the examples in Figure 3.

3.4 Official Challenge Validation

Table 6 reports the official validation results for the evaluated systems. The final direct model reaches SSIM 0.810923, MSE 0.007373, and PSNR 22.6507. The $\lambda = 0.01$ model reaches SSIM 0.811, MSE 0.007, and PSNR 22.736; the leaderboard exposes only rounded precision for this row. On the held-out fold, the same model improves SegDice from 64.22% to 65.80% overall ($p = 1.5\times10^{-10}$) and by 3.95% on gated cases ($p = 4.0\times10^{-10}$).

Table 6. Official BraTS validation results. The segmentation-consistency row is reported at the precision exposed by the public leaderboard.

Model	SSIM ↑	MSE ($\times10^{-3}$) ↓	PSNR ↑
Mask-gated control	0.735	44.500	16.061
Coarse-tissue A/B/C mean	0.793	8.710	21.950
Final direct model	**0.811**	7.373	22.651
Seg. consistency, $\lambda = 0.01$	**0.811**	**7.000**	**22.736**

3.5 Qualitative Effects of Anatomy Guidance

Figure 3 examines outputs from the final $\lambda = 0.01$ model and its control. In case (a), segmentation consistency raises 32-class SegDice from 59.5% to 70.9% and SSIM from 0.740 to 0.752. In case (b), SegDice rises from 71.7% to 76.8% while SSIM remains stable at 0.828 and 0.827. The tissue maps show that segmentation consistency changes the anatomical assignment of synthesized voxels without requiring a large intensity change. These examples illustrate why structural agreement can improve even when SSIM changes little.

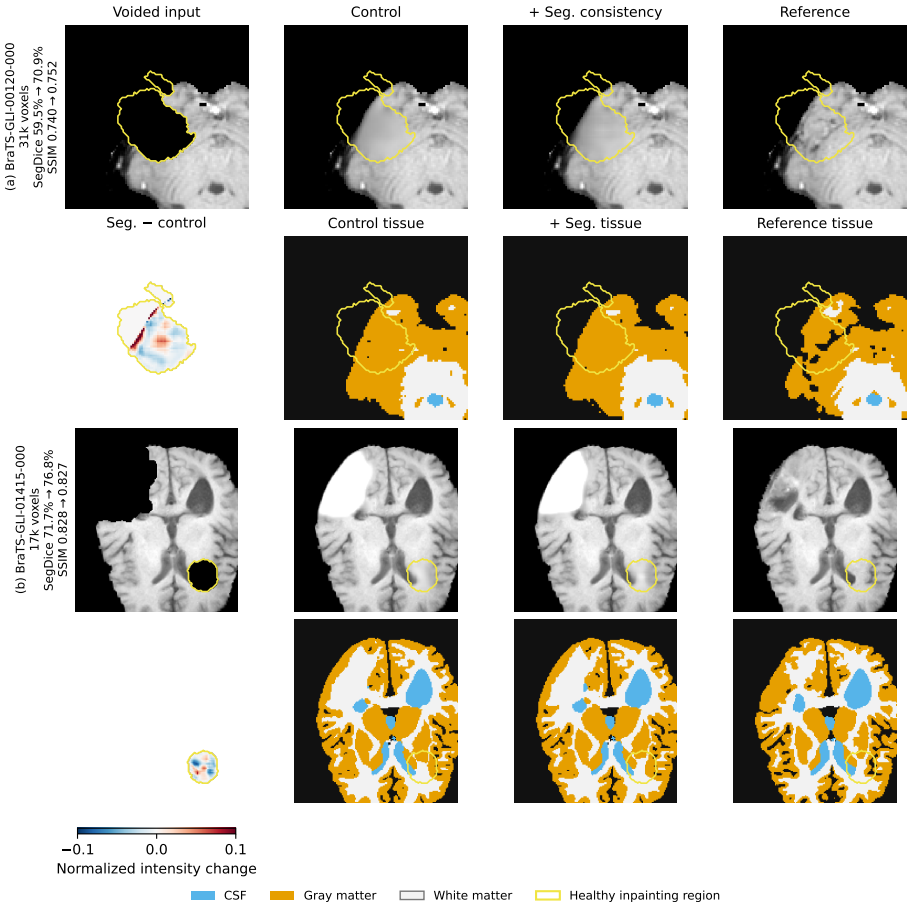

Fig. 3. Effects of the final $\lambda = 0.01$ segmentation-consistency model. Yellow contours mark the healthy inpainting region. The difference panels show normalized intensity change for segmentation consistency minus control (blue: lower; red: higher). Case (a) gains 11.4% SegDice and 0.012 SSIM; case (b) gains 5.1% SegDice while SSIM changes by less than 0.001. Tissue panels show frozen-segmenter predictions collapsed from 32 labels to CSF, gray matter, and white matter.

4 Discussion

Our results separate hidden-reference fidelity from conformity to a tissue representation. Accurate structure helps the generator, auxiliary tissue learning benefits the largest voids, and frozen-segmenter consistency changes the distribution of errors even when SSIM does not improve for an individual case or objective weight. The consistency objective can favor image cues that sharpen a frozen segmenter's tissue assignments without reproducing the reference texture.

The void-size trend is consistent with the perception–distortion trade-off [2]. A large void can admit several anatomically reasonable configurations. Regression favors a central estimate that can average uncertain detail, whereas segmentation sensitivity favors coherent tissue classes. A plausible but displaced sulcus can therefore gain segmentation overlap while losing local correlation with the single reference. The measured Dice–SSIM divergence confirms that the objectives favor different properties.

The official implementation [4] clarifies the evaluation balance. For each case, PSNR is

$$\mathrm{PSNR} = 10 \log_{10} \left(\frac{R_y^2}{\mathrm{MSE}} \right), \tag{6}$$

where R_y depends only on the normalized target. Thus, within a case, PSNR is monotonic in MSE. The challenge's equal-weight rank sum [3] consequently contains two aligned intensity-distortion orderings and one SSIM ordering. These remain the primary challenge measures; the relationship simply guides interpretation.

The anatomical experiments use one automated label space for both supervision and evaluation. SegDice therefore quantifies consistency with SynthSeg, while the official metrics quantify fidelity to the hidden image. Manual labels or expert review are required to establish neuroanatomical accuracy. Surface Dice [6], boundary distance, and tissue-volume error would extend the present regional analysis [8].

Acknowledgements

This submission is associated with the BraTS Synapse team *Yangs team*. We used BraTS 2026 Challenge data (Synapse syn74274097). Computations used institutional high-performance computing resources at the University of California, San Francisco.

Source code. `https://github.com/Guhanxue/Brain_inpainting_release`

References

1. Billot, B., Greve, D.N., Puonti, O., Thielscher, A., Van Leemput, K., Fischl, B., Dalca, A.V., Iglesias, J.E.: Synthseg: Segmentation of brain mri scans of any contrast and resolution without retraining. Medical Image Analysis **86**, 102789 (2023). https://doi.org/10.1016/j.media.2023.102789
2. Blau, Y., Michaeli, T.: The perception-distortion tradeoff. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 6228–6237 (2018). https://doi.org/10.1109/CVPR.2018.00652
3. BraTS 2026 Challenge Organizers: BraTS 2026 Challenge: Task 4 local synthesis / inpainting. https://challenges.synapse.org/Challenges/DetailsPage/Task4?id=syn74274097 (2026), accessed 29 July 2026
4. BraTS Inpainting Organizers: Official BraTS inpainting evaluation package. https://github.com/BraTS-inpainting/inpainting (2023), version 0.0.12
5. Menze, B.H., et al.: The multimodal brain tumor image segmentation benchmark (BRATS). IEEE Transactions on Medical Imaging **34**(10), 1993–2024 (2015). https://doi.org/10.1109/TMI.2014.2377694
6. Nikolov, S., Blackwell, S., Zverovitch, A., Mendes, R., Livne, M., De Fauw, J., Patel, Y., Meyer, C., Askham, H., Romera-Paredes, B., et al.: Clinically applicable segmentation of head and neck anatomy for radiotherapy: Deep learning algorithm development and validation study. Journal of Medical Internet Research **23**(7), e26151 (2021). https://doi.org/10.2196/26151
7. Park, T., Liu, M.Y., Wang, T.C., Zhu, J.Y.: Semantic image synthesis with spatially-adaptive normalization. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 2337–2346 (2019)
8. Reinke, A., et al.: Understanding metric-related pitfalls in image analysis validation. Nature Methods **21**, 182–194 (2024). https://doi.org/10.1038/s41592-023-02150-0
9. Wang, Z., Bovik, A.C., Sheikh, H.R., Simoncelli, E.P.: Image quality assessment: From error visibility to structural similarity. IEEE Transactions on Image Processing **13**(4), 600–612 (2004). https://doi.org/10.1109/TIP.2003.819861
10. Yu, J., Lin, Z., Yang, J., Shen, X., Lu, X., Huang, T.S.: Free-form image inpainting with gated convolution. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 4471–4480 (2019)
11. Zhang, J., Chen, K., Weng, Y.: Synthesis of healthy tissue within tumor area via u-net. In: Brain Tumor Segmentation, and Cross-Modality Domain Adaptation for Medical Image Segmentation. Lecture Notes in Computer Science, vol. 14669, pp. 233–240. Springer, Cham (2024). https://doi.org/10.1007/978-3-031-76163-8_21
12. Zhang, J., Weng, Y., Chen, K.: Robust 3d brain mri inpainting with random masking augmentation. In: Segmentation, Classification, and Synthesis for Brain Tumors and Traumatic Brain Injuries. Lecture Notes in Computer Science, vol. 16377, pp. 102–109. Springer, Cham (2026). https://doi.org/10.1007/978-3-032-16370-7_9