MRIComp4Flow: Compression of 3D Brain MRI for Training Multi-Modal Generative Models

Lisa K. Fischer[1*], Mykhailo Riabets[1*], Daniel Rueckert[4,5], Benedikt Wiestler[2,3], Anke Meyer-Baese[6,7] and Sandeep Nagar[2,3,7]

[1]Technical University of Munich (TUM), Germany [2]Munich Center for Machine Learning (MCML)
[3]AI for Image-Guided Diagnosis and Therapy, TUM, Germany [4]Chair for AI in Healthcare and Medicine, TUM University Hospital, Munich, Germany [5]Imperial College London [6]Florida State University [7]Institute for Advanced Study, TUM (TUM-IAS), Germany *equal contribution

PROBLEM: How far can a 3D brain MRI dataset be compressed before generative MRI synthesis quality degrades?

Compression up to 20:1 reduces storage by 12.9x with no inference slowdown.

Original BraTS 2024 dataset
40.37 GB
↓
JPEG2000 20:1
↓
3.13 GB
(12.9× smaller)
↓
Fits comfortably on local NVMe storage
Reduces training I/O burden
Easier transfer and archival
No change in inference speed (~0.35s)

Motivation: Compression as an Unanswered Question for MRI Generative Models

- Multi-modal MRI is central to brain tumor assessment. [1, 4]
- Compression is common in clinical archives and PACS systems. [2, 6]
- Moderate compression is tolerated by segmentation models. [3]
- However, generative MRI synthesis must preserve anatomy, contrast, and fine image structure.

→ Generative compression robustness remains unknown.

MRIComp4Flow Workflow:

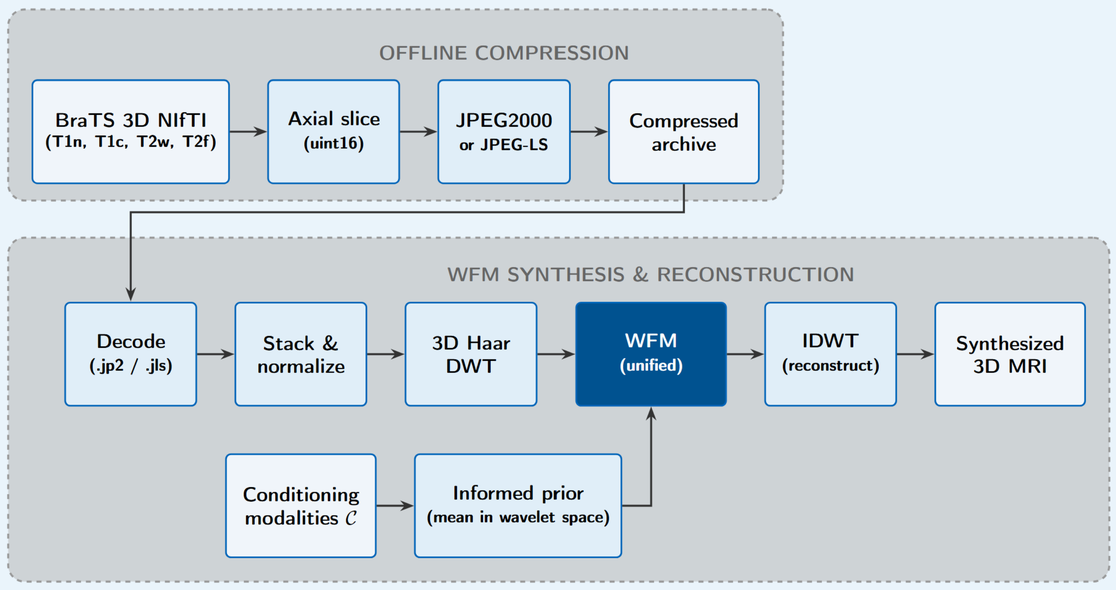

Fig. 1: MRIComp4Flow: BraTS volumes are compressed, decoded on-the-fly, and used to train a unified WFM synthesizer.

→ We train one WFM model per setting on BraTS 2024 Glioma data, evaluation included PSNR, SSIM, visual quality, and downstream tumor segmentation. [4, 5]
BraTS 2024 GLI · 82M WavUNet · 1-step inference (~0.35 second)

Our Contributions:

→ Compression pipeline: JPEG2000/JPEG-LS compression, decoding, and loading for WFM training.
→ Generative robustness analysis: novel evaluation of how MRI compression affects multi-modal generative synthesis.
→ Compression-quality trade-off: quantitative and qualitative evidence that WFM[5] remains robust at JPEG2000 compression of up to 20:1.

Qualitative Results:

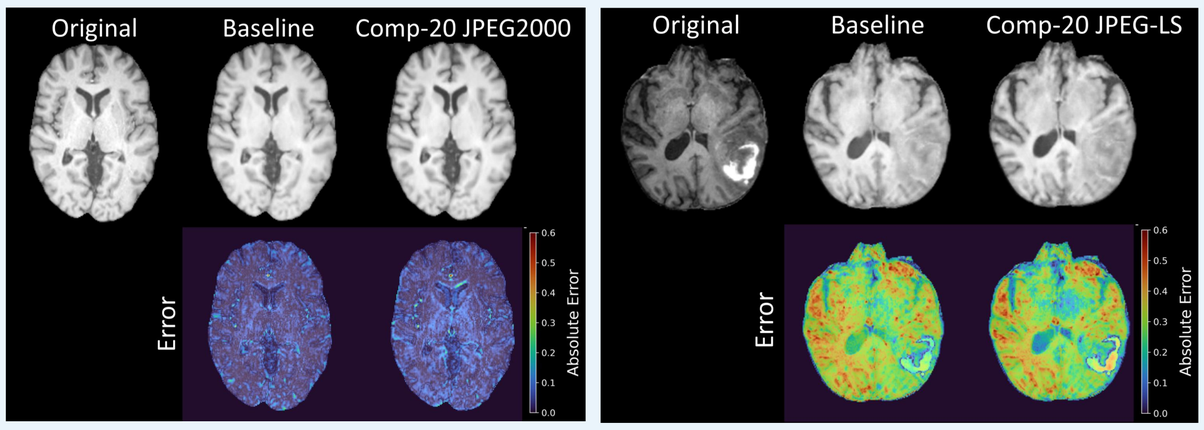

Fig. 2: Qualitative synthesis results for best (33.2 dB, left) and worst (17.3 dB, right) JPEG2000 20:1 cases, compared to the original image and baseline.

→ Tumor anatomy and contrast boundaries are largely preserved, with differences mainly confined to subtle high-frequency texture changes.

Quantitative Results:

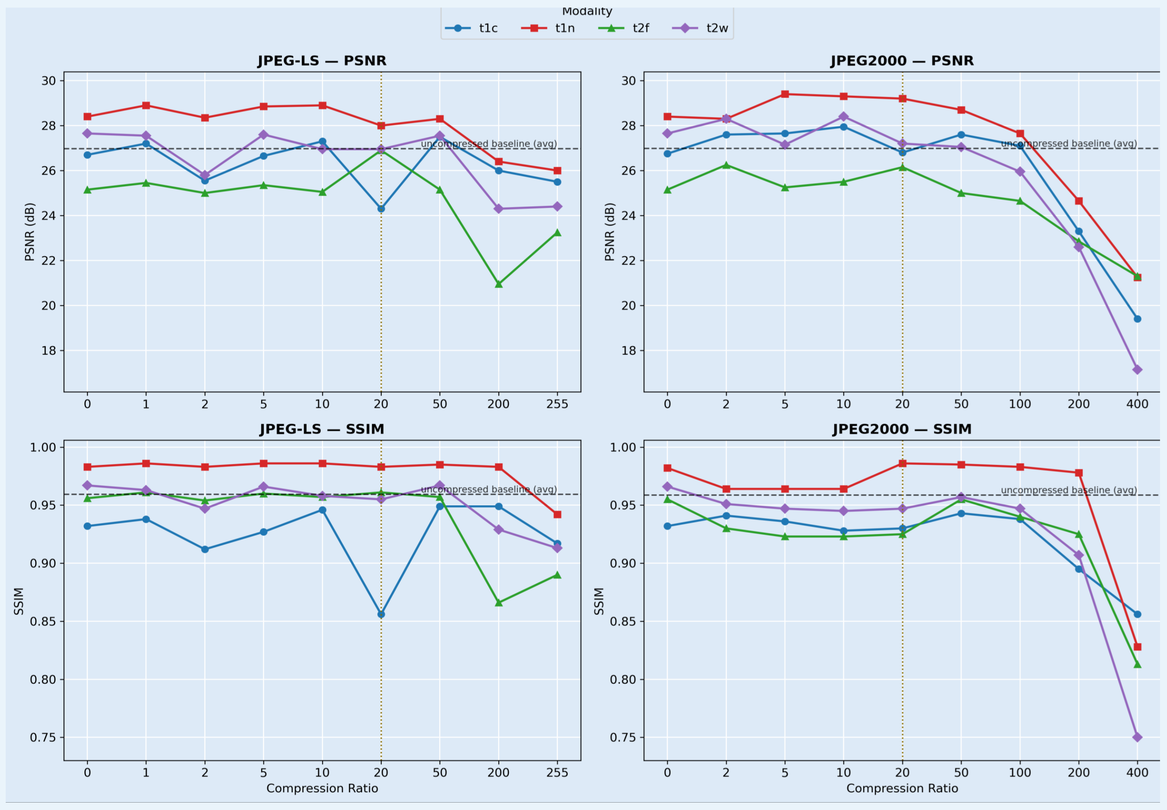

Fig. 3: PSNR and SSIM values for JPEG-LS and JPEG2000 across compression ratios and modalities.

→ Compression sometimes beats the baseline.

Tab. 1: Downstream Segmentation

	Dice WT	Dice TC
JPEG2000 2:1	0.81	0.24
JPEG2000 50:1	0.78	0.19

→ Compressed-data synthesis preserves macrostructure better than fine tumor subregions.

References:

[1] Bakas et al., BraTS Cluster of Challenges, Zenodo, 2024.
[2] Christopoulos et al., JPEG2000 Still Image Coding System, IEEE TCE, 2000.
[3] Kurmukov et al., Lossy Compression for 3D Medical Image Segmentation, arXiv 2024.
[4] Li et al. BraTS Challenge 2023: Brain MR Image Synthesis for Tumor Segmentation (BraSyn).
[5] Tur et al., WFM: 3D Wavelet Flow Matching for MRI Synthesis, MIDL 2026.
[6] Weinberger et al., JPEG-LS (LOCO-I), IEEE TIP, 2000.

Conclusion and Future Work:

→ Lossy compression up to 20:1 is a practically free lunch for 3D generative MRI: 12.9× less storage, unchanged synthesis quality and inference speed; only fine tumor substructure degrades at extreme ratios.
→ Future work could train generative models directly on compressed latent representations, exploiting the observed denoising-regularizer effect.

Acknowledgement: This work was supported by TUM – IAS Fellowship

IAS
TUM Institute for Advanced Study
Technische Universität München

AI-IDT
AI for Image-Guided Diagnosis and Therapy

Klinikum rechts der Isar
Technische Universität München

MICCAI 2026 SASHIMI Workshop

GitHub: lisafis/MRIComp4Flow