Task 8: MR Image Synthesis for BraTS

No More Slice Wars: Towards Harmonized Brain MRI Synthesis for the BraSyn Challenge

Omar Carpentiero*, Kevin Marchesini*, Costantino Grana, Federico Bolelli
University of Modena and Reggio Emilia, Italy
{name.surname}@unimore.it
*equal contribution, _ presenting the poster

Repository

1. Dataset and Preprocessing

The BraSyn-2025 dataset is derived from BraTS-GLI 2023, BraTS-METS 2023 and BraTS-Meningioma. The training set consists of 1,251 glioma cases and 238 metastasis cases. The validation set consists of 219 glioma and 31 metastasis cases. The private test set comprises 219 glioma, 59 metastasis, and 283 meningioma cases.

We standardized background voxel values across all modalities. All and only background voxels (i.e., those originally equal to zero) were reassigned a constant value of -1, which also serves as the placeholder for masked input slices.

We perform dataset-wise z-score normalization by computing the global mean (Avg.) and standard deviation (Std.) across all training volumes from both the GLI and MET datasets. These statistics are calculated after applying intensity clipping at the 99.5th percentile.

Value	Clipp.	Norm.	T1c	T1n	T2f	T2w
Max.	✗	✗	2,120,538	155,724	612,368	4,563,634
Max.	✓	✗	8,664	7,315	8,842	8,233
Avg.	✓	✗	1,066.34	781.22	510.99	673.44
Std.	✓	✗	1,301.70	944.34	769.42	804.39
Min.	✓	✓	-0.8192	-0.8273	-0.6641	-0.8372
Max.	✓	✓	5.8367	6.9189	10.8277	9.3979

2. The Architecture

Given 3 input modalities and 1 missing, the framework generates the missing modality slice-by-slice, stacking them into a full volume.

Generator:
- 4 modality-specific encoders
- Early-fusion multimodal encoder
- Channel attention block
- Transformer-based modality infuser
- Unimodal decoder

Discriminator:
- 3 convolutional blocks
- Modality classification head
- Discrimination head (fake/real classification)

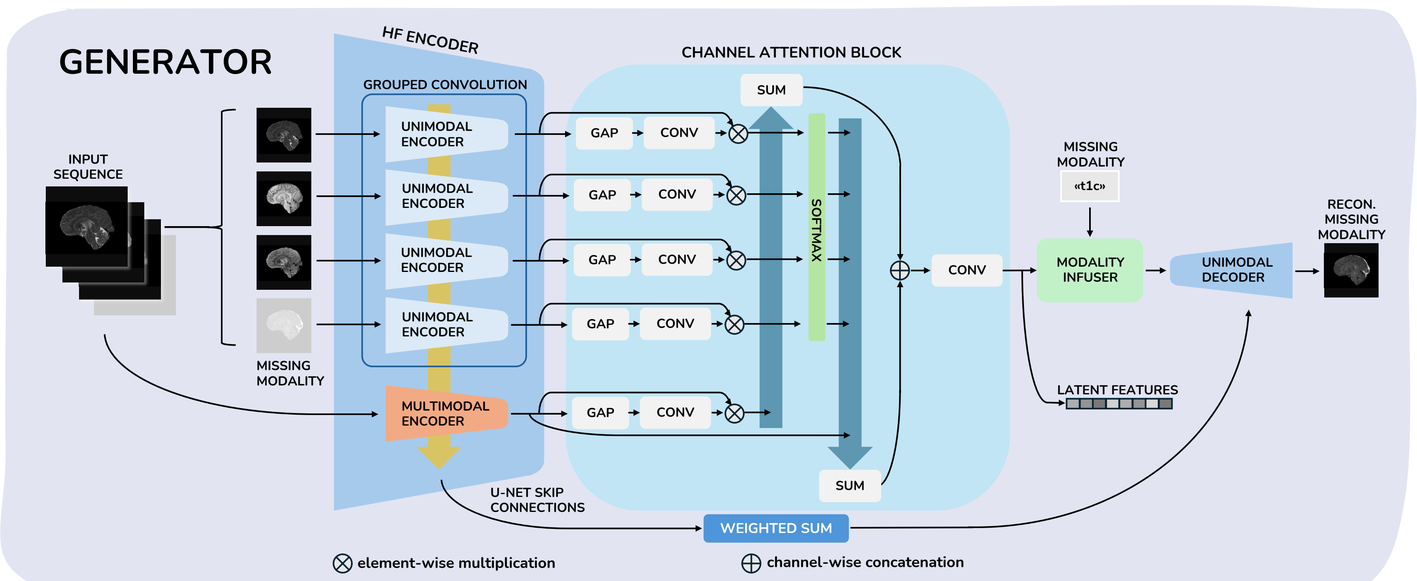

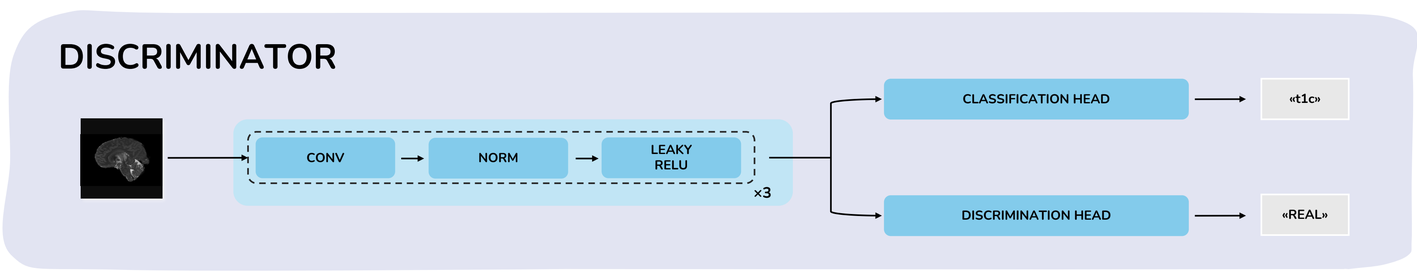

Loss:
- **Reconstruction loss**: MAE btw output and GT.
- **Adversarial loss**: penalizes the generator when the discriminator detects fake images, encouraging realism.
- **Classification loss**: enforces modality-consistent representations.
- **Cycle consistency loss**: MAE between GT and the output when a generated modality is reused as input.
- **Feature loss**: same as cycle consistency, but applied in latent feature space.
- **SSIM loss**: structural similarity output vs. GT.

LOSS

$$10 * L_{recon} + 0.25 * L_{adversarial} + 0.25 * L_{classification} + 1 * L_{feature} + 1 * L_{cycle} + 5 * L_{SSIM}$$

3. Results

Metrics are computed on the official validation set (219 GLI, 31 MET), with one randomly selected modality removed per sample. The BraTS package is used to generate pseudo-labels from the ground-truth volumes and to segment the volumes produced by our algorithm, which are then compared. Segmentation metrics are obtained from this comparison.

SSIM						DICE									NSD								
ALL		GLI		MET		ALL			GLI			MET			ALL			GLI			MET		
WT	HT	WT	HT	WT	HT	ET	TC	WT	ET	TC	WT	ET	TC	WT	ET	TC	WT	ET	TC	WT	ET	TC	WT
99.77	93.89	99.75	94.20	99.84	91.73	57.86	77.55	74.45	57.14	76.16	72.50	62.92	87.38	88.26	57.92	68.30	56.64	56.40	66.25	52.88	68.64	82.82	83.23