Tumor-Aware Direct Preference Optimization for Longitudinal Post-Operative MRI Generation

Bogyeong Kang[1], Minjoo Lim[1], Hyeonyeong Nam[1], Keun-Soo Heo[1], Mingxia Liu[2⋆], and Tae-Eui Kam[1⋆]

[1] Department of Artificial Intelligence, Korea University, Seoul, South Korea.
{kangbk, kamte}@korea.ac.kr
[2] Department of Radiology and BRIC, University of North Carolina at Chapel Hill, Chapel Hill, North Carolina 27599, USA
mingxia_liu@med.unc.edu

Abstract. Glioblastoma is a highly malignant brain tumor marked by rapid progression and high mortality, making patient-specific treatment strategies essential for improving prognosis. Longitudinal post-operative magnetic resonance imaging (MRI) plays a key role in monitoring surgical outcomes and planning subsequent therapy. Despite significant advances in generative models, longitudinal post-operative MRI generation remains challenging, as it requires reflecting *clinical preferences* that align with patient-specific post-operative outcomes. Such clinically preferred images must accurately reflect long-term post-operative trajectories that depend on the extent of resection (determined by tumor location and infiltration) and progressively evolve across follow-up time points. However, existing generative models often fail to capture these clinical preferences, leading to misalignment between generated images and clinically meaningful outcomes. To address these challenges, we propose **TADPO**, a novel Tumor-Aware Direct Preference Optimization (DPO) framework that aligns longitudinal post-operative MRI generation with clinical preferences using pre-operative MRI. Specifically, TADPO leverages DPO to generate post-operative MRIs *aligned with preferred outcomes* across multiple follow-up time points. We also introduce a *tumor-aware optimization* strategy that prioritizes tumor regions and their boundaries, improving anatomical fidelity of post-operative tumor representations and enhancing alignment with clinically meaningful preferences. TADPO demonstrates superior alignment with clinical preferences compared to state-of-the-art generative methods, highlighting its potential to support personalized treatment planning.

Keywords: MRI Generation · Direct Preference Optimization · Brain Tumor · Diffusion Model

1 Introduction

Glioblastoma is a rapidly progressing brain tumor associated with a median overall survival of under 15 months from the time of diagnosis [1, 13, 14]. While

⋆Co-corresponding authors.

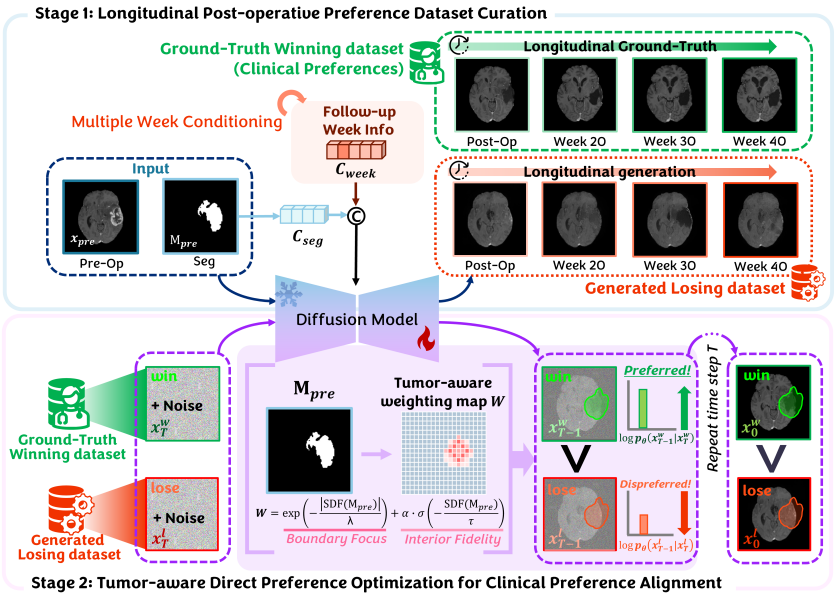

Fig. 1. TADPO for longitudinal post-operative MRI generation consists of (**Stage 1**) preference dataset curation by pairing ground-truth ("winning") and generated ("losing") MRIs and (**Stage 2**) tumor-aware DPO focused on key tumor progression regions.

standard treatment includes resection, radiotherapy, and chemotherapy, surgical outcomes significantly dictate patient prognosis [12, 21, 25]. Given that optimal treatment strategies vary across patients depending on tumor characteristics, simulating longitudinal post-operative magnetic resonance imaging (MRI) before therapy can support clinical decision-making [10]. These simulations allow clinicians to assess structural brain changes and tumor progression over time while facilitating clearer communication of expected surgical outcomes with patients [8, 20].

Unlike conventional generation tasks [9, 11, 15], longitudinal post-operative MRI generation must account for *clinical preferences* that align with patient-specific post-operative outcomes. In particular, long-term post-operative outcomes depend on the extent of resection (determined by tumor location and infiltration), and these clinically relevant factors should be reflected in generated MRIs [3, 5, 19]. Furthermore, tumor progression across follow-up time points introduces additional complexity, as accurately modeling temporal changes requires careful consideration of clinical relevance [26]. Traditional generative models struggle to incorporate such clinical preferences, often yielding images with indistinct tumor resection regions that deviate from realistic clinical outcomes.

To address these challenges, we propose **TADPO**, a novel Tumor-Aware Direct Preference Optimization (DPO) framework to generate longitudinal post-operative MRIs aligned with clinical preferences based on pre-operative MRI. As

illustrated in Fig. 1, we leverage DPO [23] to align generated outputs with clinically preferred representations at specific follow-up time points. Specifically, we curate a longitudinal post-operative preference dataset by pairing ground-truth follow-up MRIs that reflect clinically relevant factors as "winning" (preferred) images with plausibly generated but clinically inconsistent MRIs as "losing" (dispreferred) images, and optimize the model to encourage alignment with preferred outputs while avoiding dispreferred ones. Furthermore, we introduce a tumor-aware optimization strategy that prioritizes tumor regions and their boundaries, which are critical for capturing longitudinal post-operative dynamics and enhancing clinical preference alignment. Unlike prior approaches that optimize preference over the entire image [23], TADPO ensures alignment with clinically meaningful preferences in reflecting long-term tumor progression. We validate TADPO on a public dataset, highlighting its potential to support personalized treatment planning and improve prognosis.

2 Proposed Method

Preliminaries. DPO [17] is a fine-tuning method for aligning large language models to human preferences without relying on an explicit reward function. Given a dataset $\mathcal{D}$ of preference pairs (x^w, x^l, c) where x^w denotes the "winning" (preferred) samples, x^l denotes the "losing" (dispreferred) samples, and c is the conditioning input, DPO directly optimizes the model $p_\theta(x \mid c)$ by maximizing the likelihood of preferred responses over dispreferred ones while not deviating too far from the reference model $p_{\text{ref}}(x \mid c)$, as follows:

$$\mathcal{L}_{\text{DPO}}(\theta) = -\mathbb{E}_{x^w, x^l, c}\left[\log \sigma\left(\beta \log \frac{p_\theta(x^w \mid c)}{p_{\text{ref}}(x^w \mid c)} - \beta \log \frac{p_\theta(x^l \mid c)}{p_{\text{ref}}(x^l \mid c)}\right)\right], \tag{1}$$

where σ denotes the sigmoid function and β is a hyperparameter that controls the divergence between $p_\theta(x \mid c)$ and $p_{\text{ref}}(x \mid c)$.

Diffusion-DPO [23] is proposed by reformulating Eq. (1) as an objective for aligning diffusion-based generative models with human preferences. The diffusion-DPO objective is KL-regularized and aligns the reverse process with a reference model p_{ref} while maximizing human preference rewards at diffusion time step t:

$$\begin{aligned}\mathcal{L}_{\text{Diffusion-DPO}}(\theta) = -\mathbb{E}_{t, x_t^w, x_t^l}\Big[\log \sigma\Big(-\beta T\, \omega(\lambda_t)\Big(\\ \|\epsilon^w - \epsilon_\theta(x_t^w, t)\|_2^2 - \|\epsilon^w - \epsilon_{\text{ref}}(x_t^w, t)\|_2^2 \\ -\Big(\|\epsilon^l - \epsilon_\theta(x_t^l, t)\|_2^2 - \|\epsilon^l - \epsilon_{\text{ref}}(x_t^l, t)\|_2^2\Big)\Big)\Big)\Big],\end{aligned} \tag{2}$$

where $(x_0^w, x_0^l) \sim \mathcal{D}$, $t \sim \mathcal{U}(0, T)$, $x_t^w \sim q(x_t^w \mid x_0^w)$, $x_t^l \sim q(x_t^l \mid x_0^l)$, $x_t^* = \alpha_t x_0^* + \sigma_t \epsilon^*$ with $\epsilon^* \sim \mathcal{N}(0, I)$, $\lambda_t = \alpha_t^2/\sigma_t^2$ denotes the signal-to-noise ratio, and $\omega(\lambda_t)$ is a weighting function.

Proposed TADPO. Given a pre-operative MRI, our goal is to generate a post-operative MRI at a specific future follow-up time point that is aligned

with clinical preference. To achieve this, TADPO is designed as a two-stage framework: **Stage 1** curates a longitudinal post-operative preference dataset without human annotations by leveraging a pretrained reference model that has learned general knowledge of longitudinal post-operative MRI generation, and **Stage 2** further optimizes the model to generate clinically preferred outcomes by focusing on tumor regions at target follow-up time points.

Stage 1: Longitudinal Post-operative Preference Data Curation. Most Diffusion-DPO methods rely on large-scale human-aligned preference datasets [4, 23], which are unavailable in the medical domain due to the need for expert knowledge and high annotation costs. To address this, we curate a longitudinal preference dataset without requiring human annotations. Given the entire training dataset $\mathcal{D}_{\text{entire}}$, we sample a patient-level subset $\mathcal{D}_{\text{ref}} \subset \mathcal{D}_{\text{entire}}$ for pretraining the reference model p_{ref}, and use the remaining data $\mathcal{D}_{\text{dpo}} = \mathcal{D}_{\text{entire}} \backslash \mathcal{D}_{\text{ref}}$ in Stage 2. The reference model p_{ref} serves as a fixed baseline policy for the subsequent DPO stage and is additionally used to curate preference pairs. Specifically, the reference model p_{ref} is pretrained to generate post-operative MRIs at specified follow-up time points based on the pre-operative MRI x_{pre} and its segmentation mask M_{pre}, while incorporating spatial and temporal conditioning prompts (c_{seg}, c_{week}), derived from the mask and the target follow-up time point, respectively. After pretraining, for each sample in $\mathcal{D}_{\text{dpo}}$, we perform inference using the reference model p_{ref} to generate follow-up MRIs at the target time point by exploiting the model's inherent stochasticity under different conditioning prompts.

Although the images generated by the reference model p_{ref} appear visually plausible, they may remain misaligned with clinically realistic tumor progression. Thus, we treat the generated image from the reference model as the "losing" (dispreferred) sample $x_0^l \sim p_{\text{ref}}(\cdot \mid x_{pre}, c_{seg}, c_{week})$ and the corresponding ground-truth follow-up MRI as the "winning" (preferred) sample x_0^w. By pairing the losing and winning samples at each follow-up time point, we curate a longitudinal preference dataset without human annotation while maintaining strong clinical relevance. Since outputs of the reference model are used as losing samples, DPO allows the model p_θ to overcome the lack of clinical alignment inherent in the task-specific pretraining stage. Consequently, TADPO encourages the model p_θ to internalize underlying clinical progression dynamics rather than simply synthesizing plausible characteristics of the resection region, a process further stabilized by the tumor-aware weighting in subsequent optimization.

Stage 2: Tumor-Aware DPO for Clinical Preference Alignment. In traditional Diffusion-DPO [23], optimization is performed over the entire image, encouraging the generated output to align with the preferred image at a global level. However, clinically meaningful preference is primarily associated with tumor regions rather than other brain tissue structures in the entire image. In particular, tumor boundary regions exhibit the most notable temporal changes over time and are strongly associated with patient prognosis. Therefore, focusing on tumor boundary regions is especially important for capturing longitudinal post-operative dynamics. To this end, we design a novel tumor-aware DPO strategy that leverages a Signed Distance Function (SDF) to construct

a tumor-aware weighting map emphasizing both tumor boundaries and tumor core regions during preference optimization. Given a pre-operative segmentation mask M_{pre}, we define $\mathrm{SDF}(\mathrm{M}_{pre})$ such that $\mathrm{SDF}(\mathbf{p}) < 0$ for pixels $\mathbf{p}$ inside M_{pre}, and $\mathrm{SDF}(\mathbf{p}) > 0$ for pixels outside. The weighting map W is defined as follows:

$$W = \exp\left(-\frac{|\mathrm{SDF}(\mathrm{M}_{pre})|}{\lambda}\right) + \alpha \cdot \sigma\left(-\frac{\mathrm{SDF}(\mathrm{M}_{pre})}{\tau}\right), \tag{3}$$

where $\mathrm{SDF}(\mathrm{M}_{pre})$ denotes the signed Euclidean distance from M_{pre} to the tumor boundary, λ is the scale parameter controlling the decay of the boundary term, τ is the temperature parameter controlling the smoothness of the sigmoid transition, and α is the weighting parameter. This formulation consists of two specialized components: (a) **Boundary Focus:** The exponential term prioritizes the tumor boundaries (where SDF ≈ 0), ensuring the model captures high-frequency temporal changes and infiltration patterns. (b) **Interior Fidelity:** The sigmoid term (σ) assigns high weight to the tumor interior (negative SDF values), guiding the model to accurately reconstruct the resection cavity. To ensure training stability, we normalize the weighting map by its mean value. The TADPO optimizes the following objective function:

$$\begin{aligned}\mathcal{L}_{\mathrm{TADPO}}(\theta) = -\mathbb{E}_{t,x_t^w,x_t^l}\Big[\log\sigma\Big(-\beta T\,\omega(\lambda_t)\,\Big(\\ \|W \odot (\epsilon^w - \epsilon_\theta(x_t^w, t))\|_2^2 - \|W \odot (\epsilon^w - \epsilon_{\mathrm{ref}}(x_t^w, t))\|_2^2 \\ -\Big(\|W \odot (\epsilon^l - \epsilon_\theta(x_t^l, t))\|_2^2 - \|W \odot (\epsilon^l - \epsilon_{\mathrm{ref}}(x_t^l, t))\|_2^2\Big)\Big)\Big)\Big].\end{aligned} \tag{4}$$

3 Experiment

Dataset and Implementation Details. We conduct experiments on the LUMIERE [22] dataset with contrast-enhanced T1-weighted MRIs sliced into 2D axial tumor-containing images. The dataset is randomly split on a patient-level basis into training, validation, and test sets with a ratio of 8:1:1, resulting in 13,095 training images, 752 validation images, and 901 test images. The follow-up MRI scans are grouped into the post-operative time point and intervals of 1–20, 21–30, 31–40, 41–50, 51–60, 61–70, and 71–100 weeks, following standard clinical follow-up practice, with each interval represented by its upper bound [1]. TADPO is implemented based on a pretrained Stable Diffusion [18] image inpainting model. The week information is encoded as an 8-dimensional one-hot vector and embedded through a learnable layer, while the segmentation mask is processed with a strided convolution layer. Both are then projected into 1024-dimensional latent representations and concatenated as c_{week} and c_{seg} to form the conditioning prompt. For training, the dataset $\mathcal{D}_{\mathrm{entire}}$ is randomly split, with 50% assigned to $\mathcal{D}_{\mathrm{ref}}$ and the remaining 50% to $\mathcal{D}_{\mathrm{dpo}}$. For Stage 2, we set the weighting map parameters to $\lambda = 5.0$, $\tau = 8.0$, and $\alpha = 0.5$ to balance boundary sensitivity and interior reconstruction, and set $\beta = 250$ for the DPO objective.

Table 1. Quantitative evaluation of longitudinal post-operative MRI generation results against state-of-the-art diffusion-based methods. Results are reported as mean_{std}.

Method	Whole Image Evaluation				Tumor Region Evaluation			
	SSIM ↑	FSIM ↑	PSNR ↑	LPIPS ↓	SSIM ↑	FSIM ↑	PSNR ↑	LPIPS ↓
DDPM [6]	$0.655_{0.047}$	$0.844_{0.027}$	$31.92_{0.474}$	$0.263_{0.024}$	$0.581_{0.048}$	$0.879_{0.020}$	$31.63_{0.317}$	$0.261_{0.025}$
PBE [27]	$0.719_{0.054}$	$0.863_{0.023}$	$33.11_{0.757}$	$0.223_{0.027}$	$0.678_{0.048}$	$0.882_{0.018}$	$32.59_{0.627}$	$0.228_{0.026}$
InsP2P [2]	$0.684_{0.053}$	$0.852_{0.021}$	$32.56_{0.700}$	$0.224_{0.019}$	$0.678_{0.039}$	$0.883_{0.017}$	$32.17_{0.606}$	$0.225_{0.019}$
VISII [16]	$0.704_{0.036}$	$0.856_{0.019}$	$32.77_{0.745}$	$0.218_{0.018}$	$0.699_{0.033}$	$0.887_{0.016}$	$32.38_{0.659}$	$0.217_{0.018}$
ImBrush [28]	$0.715_{0.034}$	$0.866_{0.015}$	$32.48_{0.707}$	$0.220_{0.019}$	$0.680_{0.038}$	$0.887_{0.017}$	$32.05_{0.676}$	$0.219_{0.019}$
AnySD [29]	$0.689_{0.045}$	$0.865_{0.027}$	$32.52_{0.739}$	$0.224_{0.024}$	$0.666_{0.043}$	$0.884_{0.021}$	$32.19_{0.639}$	$0.222_{0.021}$
TADPO	$\mathbf{0.762}_{0.039}$	$\mathbf{0.896}_{0.024}$	$\mathbf{33.14}_{0.886}$	$\mathbf{0.202}_{0.025}$	$\mathbf{0.778}_{0.040}$	$\mathbf{0.940}_{0.013}$	$\mathbf{34.02}_{0.731}$	$\mathbf{0.162}_{0.022}$

Experimental Setup. We evaluate image generation performance using SSIM [24], FSIM [30], PSNR [7], and LPIPS [31] at both the whole image and tumor region levels, reporting the mean and standard deviation (std) for each metric and comparing TADPO with state-of-the-art methods [2,6,16,27–29]. For methods with prompts (*i.e.*, InsP2P [2], VISII [16], ImBrush [28] and AnySD [29]), we provide the text prompt "Change the MRI to a {*follow-up week info*} glioblastoma scan". DDPM [6] embeds week information for training, while PBE [27] and AnySD [29] use randomly sampled reference images from the training dataset.

Quantitative Results. To evaluate longitudinal post-operative MRI generation performance of TADPO, we compare it with established diffusion-based image generation methods [2,6,16,27–29]. As shown in Table 1, TADPO achieves superior performance across all evaluation metrics compared to competing methods. Notably, while other methods exhibit significantly degraded performance under tumor region evaluation compared to whole image evaluation, TADPO demonstrates substantially improved performance under tumor region evaluation. These results suggest that TADPO achieves improved generation performance by aligning its outputs with clinically preferred tumor region representations across longitudinal time points, resulting in long-term post-operative MRIs that exhibit structural fidelity and perceptual similarity to real images.

Qualitative Results. Fig. 2 presents qualitative comparisons of post-operative and longitudinal follow-up MRI generation across multiple time points (Week-20 to Week-70). TADPO generates temporally consistent longitudinal post-operative MRIs, whereas competing methods struggle to reflect realistic longitudinal tumor progression. Competing methods often fail to accurately reflect the tumor resection region, exhibiting distorted boundaries or even unrealistic reconstruction of normal brain tissue. In contrast, TADPO generates longitudinal post-operative MRIs that faithfully reflect the clinical preference of tumor progression. This discrepancy is further highlighted in the corresponding difference maps, where competing methods show substantial deviations within tumor regions, while TADPO demonstrates closer structural consistency with ground-truth images. These results highlight that TADPO effectively captures longitudinal tumor progression and generates clinically aligned follow-up images.

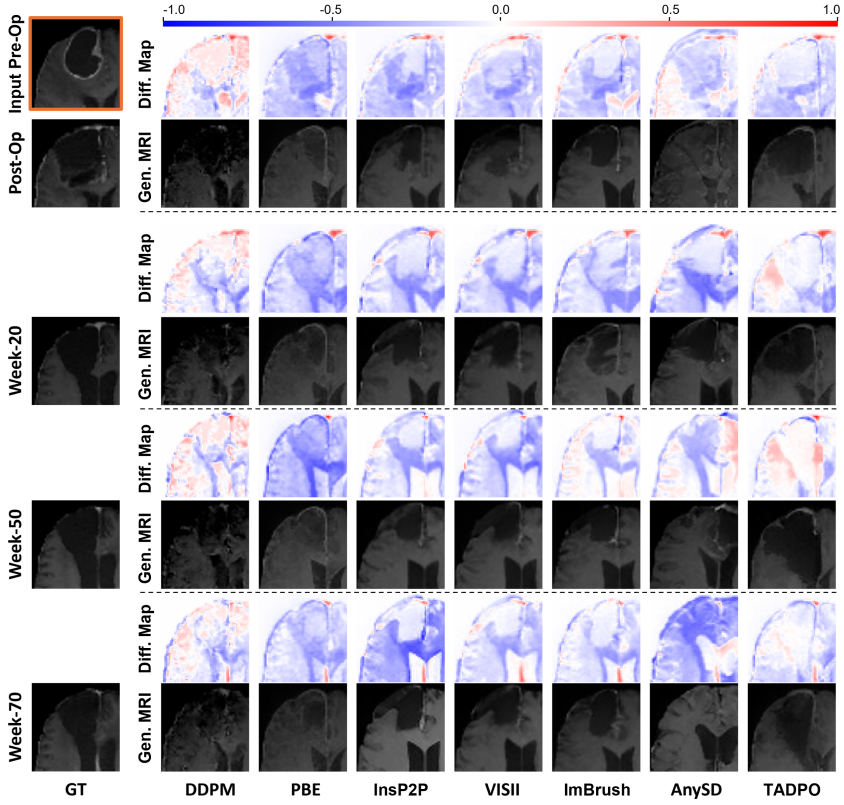

Fig. 2. Visualization of generated (Gen.) post-operative (Post-Op) and longitudinal follow-up (Week-20 to Week-70) MRIs from a pre-operative (Pre-Op) MRI. Difference (Diff.) maps between Gen. and ground-truth (GT) MRIs highlight tumor-region errors.

Table 2. Ablation study on (i) the effectiveness of tumor-aware DPO (TADPOw/oD: without preference optimization, TADPO-V: vanilla DPO [23]) and on (ii) the effectiveness of spatial (S) and temporal (T) prompts. Results are reported as mean_{std}.

Method	Whole Image Evaluation				Tumor Region Evaluation			
	SSIM ↑	FSIM ↑	PSNR ↑	LPIPS ↓	SSIM ↑	FSIM ↑	PSNR ↑	LPIPS ↓
TADPOw/oD	$0.725_{0.043}$	$0.852_{0.016}$	$33.10_{0.793}$	$0.213_{0.021}$	$0.703_{0.040}$	$0.889_{0.015}$	$32.82_{0.617}$	$0.213_{0.021}$
TADPO-V	$0.745_{0.333}$	$0.867_{0.016}$	$33.11_{0.887}$	$0.207_{0.020}$	$0.734_{0.040}$	$0.906_{0.018}$	$33.20_{0.649}$	$0.194_{0.022}$
TADPOw/oS	$0.751_{0.038}$	$0.875_{0.019}$	$33.13_{0.785}$	$0.204_{0.021}$	$0.762_{0.038}$	$0.931_{0.015}$	$33.67_{0.634}$	$0.171_{0.022}$
TADPOw/oT	$0.754_{0.033}$	$0.876_{0.019}$	$33.14_{0.889}$	$0.203_{0.022}$	$0.761_{0.038}$	$0.930_{0.015}$	$33.64_{0.629}$	$0.172_{0.022}$
TADPO	$\mathbf{0.762}_{0.039}$	$\mathbf{0.896}_{0.024}$	$\mathbf{33.14}_{0.886}$	$\mathbf{0.202}_{0.025}$	$\mathbf{0.778}_{0.040}$	$\mathbf{0.940}_{0.013}$	$\mathbf{34.02}_{0.731}$	$\mathbf{0.162}_{0.022}$

Effectiveness of Tumor-Aware DPO. To evaluate the effectiveness of the proposed tumor-aware DPO strategy, we conduct an ablation study comparing three variants. **TADPOw/oD** removes direct preference optimization, **TADPO-V** applies vanilla DPO [23] over the entire image, and TADPO denotes our proposed tumor-aware DPO strategy. Table 2 suggests that, compared to TAD-

Table 3. Quantitative analysis of the impact of data partition ratio ($\mathcal{D}_{\text{ref}}/\mathcal{D}_{\text{dpo}}$) for reference model pretraining and DPO fine-tuning. Results are reported as mean_{std}.

Ratio	Whole Image Evaluation				Tumor Region Evaluation			
	SSIM ↑	FSIM ↑	PSNR ↑	LPIPS ↓	SSIM ↑	FSIM ↑	PSNR ↑	LPIPS ↓
30/70	$0.735_{0.043}$	$0.862_{0.210}$	$33.11_{0.822}$	$0.210_{0.221}$	$0.733_{0.039}$	$0.906_{0.019}$	$33.22_{0.658}$	$0.196_{0.211}$
50/50	$\mathbf{0.762}_{0.039}$	$\mathbf{0.896}_{0.024}$	$\mathbf{33.14}_{0.886}$	$\mathbf{0.202}_{0.025}$	$\mathbf{0.778}_{0.040}$	$\mathbf{0.940}_{0.013}$	$\mathbf{34.02}_{0.731}$	$\mathbf{0.162}_{0.022}$
70/30	$0.740_{0.034}$	$0.868_{0.019}$	$33.11_{0.809}$	$0.209_{0.022}$	$0.749_{0.038}$	$0.907_{0.018}$	$33.51_{0.631}$	$0.195_{0.021}$

POw/oD, both TADPO-V and TADPO achieve improved performance, indicating that preference alignment through DPO is important for longitudinal post-operative MRI generation. TADPO-V exhibits inferior performance compared to TADPO, particularly in tumor-region evaluation. This implies that global image preference optimization does not sufficiently capture clinically relevant preferences in tumor regions, which are critical for generating longitudinal images. This confirms that TADPO's tumor-aware strategy better reflects clinical preferences, making it highly effective for longitudinal MRI generation.

Effectiveness of Spatial and Temporal Prompts. An ablation study is conducted to assess the impact of conditioning prompts in the diffusion model. Specifically, **TADPOw/oS** removes the segmentation mask from the prompt, while **TADPOw/oT** removes the follow-up time point information. Table 2 shows that TADPOw/oS exhibits a significant drop in performance compared to TADPO. This suggests that the segmentation mask provides spatial information about tumor regions, and its absence reduces the ability to capture tumor location and structure. TADPOw/oT underperforms relative to TADPO, as it struggles to capture temporal variations in tumor progression and distinguish longitudinal differences across follow-up weeks. This highlights the importance of spatial and temporal prompts in longitudinal post-operative MRI generation.

Impact of Data Partition for Pretraining and DPO. We further study the impact of different data partition ratios between reference model pretraining set $\mathcal{D}_{\text{ref}}$ and DPO set $\mathcal{D}_{\text{dpo}}$. In Table 3, the ratio "a/b" indicates that a% of the training data is allocated to pretraining $\mathcal{D}_{\text{ref}}$, and the remaining b% is used for DPO preference curation and fine-tuning $\mathcal{D}_{\text{dpo}}$. TADPO with a ratio of 30/70 underperforms because limited pretraining data prevents the reference model from acquiring sufficient task-specific knowledge, and subsequent DPO fine-tuning cannot fully compensate for this deficiency. With a 70/30 ratio, we achieve slightly better results, suggesting that increased pretraining data improves task-specific prior learning. However, its performance is still limited because the smaller DPO subset $\mathcal{D}_{\text{dpo}}$ provides fewer curated losing samples, which reduces the effectiveness of clinical preference alignment. TADPO with a ratio of 50/50 achieves the best performance, indicating that a balanced split provides sufficient data for both prior learning and clinical preference alignment.

4 Conclusion and Future Work

We propose TADPO, a novel framework for longitudinal post-operative MRI generation with clinical preference alignment. By leveraging a tumor-aware DPO strategy, TADPO focuses on tumor regions to model longitudinal tumor progression and improve clinical alignment at specific follow-up time points. Extensive experiments demonstrate the effectiveness of TADPO in longitudinal post-operative MRI generation and its potential to support treatment planning. Future work will focus on extending TADPO to multi-modal and multi-dataset settings and performing DPO in the voxel space based on 3D MRI data.

Acknowledgments. This work was supported by grants from the Institute of Information & Communications Technology Planning & Evaluation (IITP), specifically the Artificial Intelligence Graduate School Program at Korea University (No. RS-2019-II190079), the National Research Foundation of Korea (NRF) funded by the Korea government (MSIT) (Nos. RS-2023-00212498 and RS-2025-25302986), and the Korean ARPA-H Project through the Korea Health Industry Development Institute (KHIDI), funded by the Ministry of Health & Welfare, Republic of Korea (No. RS-2025-25455839).

Disclosure of Interests. The authors have no competing interests to declare that are relevant to the content of this article.

References

1. Booth, T.C., Luis, A., Brazil, L., Thompson, G., Daniel, R.A., Shuaib, H., Ashkan, K., Pandey, A.: Glioblastoma post-operative imaging in neuro-oncology: Current UK practice (GIN CUP study). European Radiology **31**, 2933–2943 (2021)
2. Brooks, T., Holynski, A., Efros, A.A.: InstructPix2Pix: Learning to follow image editing instructions. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 18392–18402 (2023)
3. Brown, T.J., Brennan, M.C., Li, M., Church, E.W., Brandmeir, N.J., Rakszawski, K.L., Patel, A.S., Rizk, E.B., Suki, D., Sawaya, R., et al.: Association of the extent of resection with survival in glioblastoma: A systematic review and meta-analysis. JAMA Oncology **2**(11), 1460–1469 (2016)
4. Croitoru, F.A., Hondru, V., Ionescu, R.T., Sebe, N., Shah, M.: Curriculum direct preference optimization for diffusion and consistency models. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 2824–2834 (2025)
5. de Haan, Y., Moussa, A., Gerritsen, J., Schnitzler, S., Dirven, C., Vincent, A.: P27. 03. b biopsy versus resection in glioblastoma treatment: A propensity score-matched analysis of a retrospective cohort. Neuro-Oncology **26**, 136–137 (2024)
6. Ho, J., Jain, A., Abbeel, P.: Denoising diffusion probabilistic models. Advances in Neural Information Processing Systems **33**, 6840–6851 (2020)
7. Hore, A., Ziou, D.: Image quality metrics: Psnr vs. ssim. In: 2010 20th International Conference on Pattern Recognition. pp. 2366–2369. IEEE (2010)

8. Huang, G., Jin, R., Tang, Y., Zhao, C., Harada, T., Li, X., Lin, G.: Interactive tumor progression modeling via sketch-based image editing. arXiv preprint arXiv:2503.06809 (2025)
9. Kang, B., Nam, H., Kang, M., Heo, K.S., Lim, M., Oh, J.H., Kam, T.E.: Target-aware cross-modality unsupervised domain adaptation for vestibular schwannoma and cochlea segmentation. Scientific Reports **14**(1), 27883 (2024)
10. Kang, B., Park, S.J., Lim, M., Kang, M., Heo, K.S., Oh, J.H., Lee, H.J., Kam, T.E.: Pre-to-post operative mri generation with retrieval-based visual in-context learning. In: International Conference on Medical Image Computing and Computer-Assisted Intervention. pp. 534–544. Springer (2025)
11. Kim, J., Park, H.: Adaptive latent diffusion model for 3d medical image to image translation: Multi-modal magnetic resonance imaging study. In: Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision. pp. 7604–7613 (2024)
12. Kowalczuk, A., Macdonald, R.L., Amidei, C., Dohrmann III, G., Erickson, R.K., Hekmatpanah, J., Krauss, S., Krishnasamy, S., Masters, G., Mullan, S.F., et al.: Quantitative imaging study of extent of surgical resection and prognosis of malignant astrocytomas. Neurosurgery **41**(5), 1028–1038 (1997)
13. Kubben, P.L., ter Meulen, K.J., Schijns, O.E., ter Laak-Poort, M.P., van Overbeeke, J.J., van Santbrink, H.: Intraoperative MRI-guided resection of glioblastoma multiforme: A systematic review. The Lancet Oncology **12**(11), 1062–1070 (2011)
14. McGirt, M.J., Villavicencio, A.T., Bulsara, K.R., Friedman, A.H.: MRI-guided stereotactic biopsy in the diagnosis of glioma: Comparison of biopsy and surgical resection specimen. Surgical Neurology **59**(4), 279–283 (2003)
15. Meng, X., Sun, K., Xu, J., He, X., Shen, D.: Multi-modal modality-masked diffusion network for brain mri synthesis with random modality missing. IEEE Transactions on Medical Imaging **43**(7), 2587–2598 (2024)
16. Nguyen, T., Li, Y., Ojha, U., Lee, Y.J.: Visual instruction inversion: Image editing via image prompting. Advances in Neural Information Processing Systems **36**, 9598–9613 (2023)
17. Rafailov, R., Sharma, A., Mitchell, E., Manning, C.D., Ermon, S., Finn, C.: Direct preference optimization: Your language model is secretly a reward model. Advances in Neural Information Processing Systems **36**, 53728–53741 (2023)
18. Rombach, R., Blattmann, A., Lorenz, D., Esser, P., Ommer, B.: High-resolution image synthesis with latent diffusion models. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 10684–10695 (2022)
19. Sanai, N., Berger, M.S.: Glioma extent of resection and its impact on patient outcome. Neurosurgery **62**(4), 753–766 (2008)
20. Soto-Gamez, A., Gunawan, J.P., Barazzuol, L., Pringle, S., Coppes, R.P.: Organoid-based personalized medicine: from tumor outcome prediction to autologous transplantation. Stem Cells **42**(6), 499–508 (2024)
21. Stupp, R., Mason, W.P., Van Den Bent, M.J., Weller, M., Fisher, B., Taphoorn, M.J., Belanger, K., Brandes, A.A., Marosi, C., Bogdahn, U., et al.: Radiotherapy plus concomitant and adjuvant temozolomide for glioblastoma. New England Journal of Medicine **352**(10), 987–996 (2005)
22. Suter, Y., Knecht, U., Valenzuela, W., Notter, M., Hewer, E., Schucht, P., Wiest, R., Reyes, M.: The LUMIERE dataset: Longitudinal Glioblastoma MRI with expert RANO evaluation. Scientific Data **9**(1), 768 (2022)
23. Wallace, B., Dang, M., Rafailov, R., Zhou, L., Lou, A., Purushwalkam, S., Ermon, S., Xiong, C., Joty, S., Naik, N.: Diffusion model alignment using direct preference

optimization. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 8228–8238 (2024)
24. 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)
25. Weller, M., Le Rhun, E., Preusser, M., Tonn, J.C., Roth, P.: How we treat glioblastoma. ESMO Open **4**, e000520 (2019)
26. Wu, L., Chai, R., Lin, Z., Wu, R., Yao, D., Jiang, T., Wang, Q.: Evolution-driven crosstalk between glioblastoma and the tumor microenvironment. Cancer Biology & Medicine (2023)
27. Yang, B., Gu, S., Zhang, B., Zhang, T., Chen, X., Sun, X., Chen, D., Wen, F.: Paint by example: Exemplar-based image editing with diffusion models. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 18381–18391 (2023)
28. Yang, Y., Peng, H., Shen, Y., Yang, Y., Hu, H., Qiu, L., Koike, H., et al.: ImageBrush: Learning visual in-context instructions for exemplar-based image manipulation. Advances in Neural Information Processing Systems **36**, 48723–48743 (2023)
29. Yu, Q., Chow, W., Yue, Z., Pan, K., Wu, Y., Wan, X., Li, J., Tang, S., Zhang, H., Zhuang, Y.: AnyEdit: Mastering unified high-quality image editing for any idea. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 26125–26135 (2025)
30. Zhang, L., Zhang, L., Mou, X., Zhang, D.: Fsim: A feature similarity index for image quality assessment. IEEE Transactions on Image Processing **20**(8), 2378–2386 (2011)
31. Zhang, R., Isola, P., Efros, A.A., Shechtman, E., Wang, O.: The unreasonable effectiveness of deep features as a perceptual metric. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 586–595 (2018)