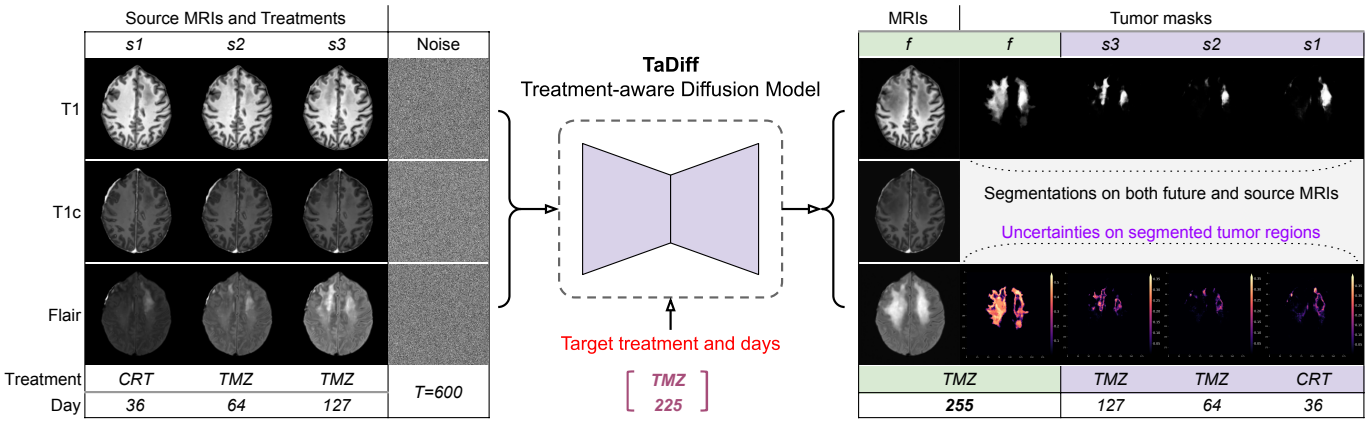

Fig. 1. An overview of the TaDiff model (short for Treatment-aware Diffusion Probabilistic model). The goal of our method is to generate a set of synthetic MRIs and tumor progression masks for any given target/future treatment (e.g., TMZ: temozolomide) and time point (e.g., Day: 225) with source sequential MRIs (e.g., *s1*, *s2*, and *s3*) and treatments (e.g., CRT: chemoradiation at Day 36, TMZ at Days 64 and 127). More details are presented in Section III.

It is noteworthy that the reverse conditional probability is tractable when conditioned on $\mathbf{x}_0$:

$$q(\mathbf{x}_{t-1}|\mathbf{x}_t, \mathbf{x}_0) = \mathcal{N}(\mathbf{x}_{t-1}; \tilde{\boldsymbol{\mu}}_t(\mathbf{x}_t, \mathbf{x}_0), \tilde{\beta}_t \mathbf{I}) \ , \tag{6}$$

where

$$\tilde{\boldsymbol{\mu}}_t \left(\mathbf{x}_t, \mathbf{x}_0\right) = \frac{\sqrt{\bar{\alpha}_{t-1}}\beta_t}{1-\bar{\alpha}_t}\mathbf{x}_0 + \frac{\sqrt{\alpha_t}\left(1-\bar{\alpha}_{t-1}\right)}{1-\bar{\alpha}_t}\mathbf{x}_t \ , \tag{7}$$

and

$$\tilde{\beta}_t = \frac{1-\bar{\alpha}_{t-1}}{1-\bar{\alpha}_t}\beta_t \ . \tag{8}$$

because of $\mathbf{x}_0 = \frac{1}{\sqrt{\bar{\alpha}_t}}(\mathbf{x}_t - \sqrt{1-\bar{\alpha}_t}\boldsymbol{\epsilon}_t)$ (Eq. 2), then

$$\tilde{\boldsymbol{\mu}}_t = \frac{1}{\sqrt{\alpha_t}}\left(\mathbf{x}_t - \frac{1-\alpha_t}{\sqrt{1-\bar{\alpha}_t}}\boldsymbol{\epsilon}_t\right) . \tag{9}$$

3) Training: For the reverse diffusion process, a neural network is trained to approximate the conditional probability distributions, i.e., train $\boldsymbol{\mu}_\theta$ to predict $\tilde{\boldsymbol{\mu}}_t$. Because $\mathbf{x}_t$ is available (Eq. 9) as input in training time, it is common to predict $\boldsymbol{\epsilon}$ from the input $\mathbf{x}_t$ at time step t, thus

$$\tilde{\boldsymbol{\mu}}_t \approx \boldsymbol{\mu}_\theta(\mathbf{x}_t, t) := \frac{1}{\sqrt{\alpha_t}}\left(\mathbf{x}_t - \frac{1-\alpha_t}{\sqrt{1-\bar{\alpha}_t}}\tilde{\boldsymbol{\epsilon}}_\theta(\mathbf{x}_t, t)\right) . \tag{10}$$

By letting $\boldsymbol{\Sigma}_\theta(\mathbf{x}_t, t) = \tilde{\beta}_t \mathbf{I}$, and letting the forward variances β_t to be a sequence of linearly increasing constants from $\beta_1 = 10^{-4}$ to $\beta_T = 0.02$, and some other simplifications in the work [26], we can minimize the MSE loss of the noise to train the neural network.

$$\mathbb{E}_{t\sim[1,T],\mathbf{x}_0,\boldsymbol{\epsilon}}\left[\|\boldsymbol{\epsilon} - \tilde{\boldsymbol{\epsilon}}_\theta(\mathbf{x}_t, t)\|^2\right] . \tag{11}$$

4) Inference: A neural network trained in the reverse diffusion process can be used to generate data. This is achieved by initializing $\mathbf{x}_T \sim \mathcal{N}(\mathbf{0}, \mathbf{1})$ and, in T steps, denoising the image by using

$$\mathbf{x}_{t-1} = \frac{1}{\sqrt{\alpha_t}}\left(\mathbf{x}_t - \frac{1-\alpha_t}{\sqrt{1-\bar{\alpha}_t}}\tilde{\boldsymbol{\epsilon}}_\theta\left(\mathbf{x}_t, t\right)\right) + \sqrt{\tilde{\beta}_t}\boldsymbol{z} \ . \tag{12}$$

where $\mathbf{z} \sim \mathcal{N}(\mathbf{0}, \mathbf{1})$ is new noise added between each denoising step.

III. Methods

The classical DDPM approach requires only $\mathbf{x}_t$ for training, resulting in arbitrary images $\mathbf{x}_0$ when sampling from random noise during inference. However, our goal is not to generate arbitrary images but to generate realistic MRIs and tumor growth maps for any target (future) treatment-day point from a given sequence of source/conditioning images and treatment information. To this end, we propose the treatment-aware diffusion (TaDiff) model for multi-parametric MRI generation and tumor growth prediction on longitudinal data. Our TaDiff model introduces a treatment-aware mechanism for conditioning a diffusion model while also employing a joint learning strategy to segment the tumor and project its future growth during diffusion processes. Figure 2 illustrates an overview of the TaDiff pipeline.

A. Problem Settings

Let tumor binary masks $\mathbf{M} \in \mathbb{R}^{L\times H\times W\times D}$ be longitudinal 3D tumor volumes with temporal length L. The corresponding longitudinal MRI scans $\mathbf{X} \in \mathbb{R}^{L\times C\times H\times W\times D}$ with C channels. In the current study, we consider $C = 3$ due to the availability of three inputs: T1-weighted (T1), contrast-enhanced T1 (T1C), and fluid-attenuated inversion recovery (FLAIR) images. The corresponding treatment information is represented as $\mathcal{T} = \{\tau_1, \tau_2, \ldots, \tau_l, \ldots, \tau_L\}$, indicating the treatment distribution, with the associated treatment days defined as $\mathcal{D} = \{d_1, d_2, \ldots, d_l, \ldots, d_L\} \quad \forall \quad d \in \mathbb{N}_0$ and $0 \leq d_{l-1} < d_l$. This work considers two treatment types: chemoradiation (CRT) and temozolomide (TMZ), specified as $\tau \in \{1, 2\} \sim \mathcal{T}$.

We randomly sample a sorted sequence of three scalar indices from available longitudinal exams as conditional sources, i.e. $\mathcal{S} = \{s_1, s_2, s_3\}$, such that $s_i \in [1, \ldots, L-1]$ and $s_i \leq s_{i+1}$. Then we sample a scalar index of future (target) sessions from the rest of future exams, that is, $f \in [s_3 + 1, \ldots, L]$. The set of conditional MRIs $\mathbf{X}$ is $\mathbf{X}^{\mathcal{S}} \in \mathbb{R}^{3\times C\times H\times W\times D}$ and the set of future/target MRIs is $\mathbf{X}^f \in \mathbb{R}^{1\times C\times H\times W\times D}$, correspondingly, we also get the